Towards spinor condensates at ultralow magnetic field

Creating dipolar quantum gases

Mark Brannan

A thesis presented for the degree of
Doctor of Philosophy

UNIVERSITY OF
BIRMINGHAM

School of Physics and Astronomy
University of Birmingham
United Kingdom

Submitted: September 20, 2016

Abstract

This thesis outlines the design and construction of an experiment to investigate dipolar interactions within a rubidium-87 (^{87}Rb) spinor Bose-Einstein condensate (BEC). For dipolar interactions to drive population transfer between the internal states of a ^{87}Rb BEC the ambient magnetic field must be reduced below 100 pT. To this end I have designed an active magnetic stabilisation system which is able to reduce the ambient magnetic field to a given level with 30 nT peak-to-peak fluctuations. A design is presented for a complementary passive magnetic shield in order to further reduce the ambient magnetic field, which electromagnetic simulations show should provide a factor of $> 10^8$ reduction in the ambient magnetic field.

To Archie and Louise

Acknowledgements

Mat Turnbull — thank you for welcoming me in and introducing me to your friends and the wider social group in Birmingham. It was great fun to have you as my partner in crime in the office, events organiser, news feed, and general fountain of internet and movie trash. Lyndsie Smith — you've been a like-minded fun-liker, fellow trouble-maker, and a promising smash-apprentice. Taking time out to talk with you was always good fun. I've enjoyed grossing you out, making you laugh, and telling you about my adventures through Wikipedia. Harry Proud — finally, someone in the lab with good taste in music. It has also helped a lot that you know what's going on with atomic physics. Discussions with you about physics and the rest of the world have always been really enjoyable. Wei He — ever positive and always willing to help whenever you can, you've been a great person to have around these past few years. The funniest quotes over the years have always come from you. Insights into your mind and the cultural differences between China and the UK have been an education. Nadine Meyer — the only person to ever scare me with how productive they can be. Having you around the office to bounce ideas and stationary off of has been both fun and helpful (for me at least). Thanks for never kicking the shit out of me even though I probably deserved it more than once. You've been something of a mother figure for the group and I appreciate all your guidance and firm hand when I needed it. Alexander Niggebaum — my fellow maker and general thing-creator. I will never know how you manage to produce so much stuff at such a high standard. You always seemed to already be great at everything. Sharing ideas and time with you has been really helpful. Vincent Boyer — thank you for taking time to discuss my thesis plan and offering advice and help when I needed it most. Chris Gill — I'm pretty sure you already know everything. Running ideas by you or just listening to yours has always been enlightening.Jon Goldwin — for sanity checking my ideas, helping debug electronics, and for graciously allowing me the

use of lab space for way longer than expected. Georgios Voulazeris — it was good to have someone else as interested in magnetic shielding as me. Our discussions about magnetic shielding were very helpful. Steve Brookes, Antony Barnett, David Taylor, Jag Sangha, Ben, Barry — thank you for all your excellent mechanical work, teaching me how to use the machines, and for just generally being a great group of guys. You have all been central to the success of our experiment and have always went the extra mile to realise all of the complex parts we designed. Chris Collins — the omniscient spectre, ghost, conscience, brick fighter, bottle soldier, and hammer club member. We've written some pretty tall tales over the past few years and all of our crazy scrapes have been amazing fun. I've long stopped trying to keep up with you and hope to someday gain or osmose your etenity's worth of knowledge. Aisling Johnson — for making me remember what life is about. Time spent sharing music, physics, time, and food with you made me the happiest I'd been for a long time. A PhD might have been worth it if only to have met you. Dominik Vogel — thoughtful, sensitive and caring. Meeting and spending time with you has been really enjoyable. I hope to be as relaxed, speak as many languages, and answer as few emails as you one day. Jamie Law — for sending expensive gift items, I will. Haaa. You were totally right about the PhD. Working alongside you prepared me well to turn work into fun. Regards, Doln. Brigid Jones — for putting up with my pish and taking care of me when I needed it. Katie McCandlish — thank you for always looking out for me. You've helped me through all of the highs and lows of the PhD. I've never been too far away for a letter to reach and I've enjoyed showing you around bustling cities when we have the chance. I'm sure there'll be plenty more to show you in the future. Alexander Rayner — for being a friend and bringing me back down to Earth when I was stuck in the PhD bubble. Mum — thanks for always having a place for me to sleep when I turn up unannounced, driving me around, and feeding me. I'm awful at staying in touch, but you never held that against me. You don't try to interfere with my life, as much as you might want to, and I really appreciate that.

Thanks for giving me the freedom to do what I wanted to: fuck things up and figure them out on my own. I'm doing a pretty good job so far. I love you. Dad — thank you for helping me make the first important steps through life, moving me up to Wick and down to Birmingham. You've literally travelled the length of the country for me. I'm really grateful for all the time, thought, and advice you've given me and hope I can use them to become a bit like you some day. You're one of the coolest guys I've ever met and I'm glad you're my dad. I love you. Daniel Benedicto Orenes — we've been through a lot (of broken experiments) over the past few years and I've not always been the easiest person to be around. Thanks for putting up with my shit and helping keep the experiment moving forward even when I felt like giving up, you've really helped a lot. Anna Kowalczyk — thanks for all your direction and guidance (sometimes literally) over the last few years, I think we've both learned a lot. Your determination to carry on even when things have seemed hopeless is admirable and has no doubt kept the experiment going. Tilman Zibold — thanks for all of the knowledge I was able to osmose during our time together. You have a way of explaining with the least possible words and math (when you want to), which helped me to pick up a working knowledge of many topics very quickly. You have been a good friend and have provided some interesting, fun conversation and antics. Fabrice Gerbier — thanks for allowing me to work in your lab alongside all of the amazing students in your group. I have used the knowledge I gained during my short stay almost every week during my PhD and I am grateful to have been given such a good grounding and introduction to atomic physics research. Kai Bongs — throughout the course of my PhD I have been able to completely immerse myself in my work with complete freedom of approach and implementation, and for this I am truly grateful. Your sage advice has always set me on the right path, or at least in the right direction. Anything was open for discussion and everything able to change, which has been very satisfying. Seeing the cold atoms group grow, receive regular substantial funding from a plethora of sources and be associated with all of

the important researchers, company partners, and industry you have made contact with has been a dizzying workload to consider, and one which you have carried out expertly; Continually adapting to your ever-changing role within the department and university. I am really proud of your achievements and what you have been able to do. If I can only become half as knowledgeable and productive as you in future then I'll be doing pretty damn well. Erling Riis — you gave me advice when I needed it most and have always been nothing but straight to the point. Your help during my final year in Strathclyde was much appreciated and has stood me in good stead until now. Iain Baikie — thank you for the advice you gave to me during my time at KP Technology, it has certainly shaped my career thus far and I still use the skills and advice I gained from working with you to this day. Andreas Freise — thanks for looking out for me and for your insights into academia. The past four years have certainly taught me that *per ardua ad alta*.

The worst is not certain
Jean Dalibard

There's the pendulum ticking, it doesn't care, it doesn't feel sorry for anyone
Fyodor Dostoyevsky

Le temps detruit tout/Time destroys all things
Gaspar Noé

Don't get any big ideas, they're not gonna happen
Thom Yorke

Happiness only real when shared
Christopher McCandless

I don't know what I'm doing but I gotta get things done
Vapours

Contents

1 Introduction **1**

2 Theory overview **7**

2.1 Atomic configuration 7

2.2 The magneto-optical trap 10

2.2.1 The Doppler shift 11

2.2.2 Optical molasses 12

2.2.3 The Zeeman effect 14

2.3 The dipole trap 18

2.3.1 The dipole force 19

2.3.2 Evaporative cooling 21
2.4 Optical lattices 22
2.5 Bose-Einstein condensation 23
2.5.1 Interactions within a condensate 24
2.5.2 Spinor BEC 26
2.5.3 Dipolar BEC 26

3 Experiment **29**
3.1 Vacuum 30
3.1.1 Pumps 33
3.1.2 Alkali source 35
3.1.3 2D MOT: Windows 37
3.1.4 Indium sealing 39
3.1.5 2D MOT: Optics 47
3.1.6 2D MOT: Magnetic fields 49
3.1.7 2D–3D MOT population transfer 54
3.1.8 HV-UHV separation 57
3.1.9 UHV chamber 58
3.1.10 Glass cell: design 59
3.1.11 Glass cell: construction 62
3.2 Optics 66
3.2.1 Locking and distribution 67
3.2.2 Dipole trap 73
3.2.3 Optical lattice 73
3.2.4 Imaging 75
3.2.5 Non-magnetic optics 77
3.3 Passive magnetic shielding 79
3.3.1 Theory overview 81
3.3.2 Simulating shielding performance 92

3.3.3 Shield design 93
3.3.4 Simulations 94

4 Active magnetic shielding **99**

4.1 Theory 101
4.2 Design 102
4.2.1 Components 103
4.2.2 Bandwidth and noise 105
4.2.3 Offset 107
4.3 Performance 107

5 Outlook **113**

References **117**

1

Introduction

Magnetism has long been easily observed but difficult to explain and understand. Qualitative studies are recorded as far back as 380 BCE by Euripedes (referenced by Socrates in a dialogue with Ion)[1], circa 150 CE by Lucretius[2], 78 CE by Pliny the Elder[3], in 1589 by Porta[4], and later in 1600 by Gilbert[5]. Early accounts focus on the natural magnetism present in lodestone, a naturally occuring magnetised form of magnetite mineral. These early encounters with magnetism are met with wonder and are described more often by poems than mathematics. Socrates relates magnetism to that of the inspiration of a man by a muse[1], Lucretious describes the allure of the effects of magnetism of the lodestone in his poem "On the Nature of Things" saying *"A stone there is that men admire much, that makes rings hang*

in chains by touch"[2], and Porta recounts his delight at his first observation of magnetic attraction in his books on "Natural Magick", saying that *"the [lodestone] will wonderfully draw iron, which when I first saw, I trembled at exceedingly"*[4].

Further work over hundreds of years led to the invention of electricity whereupon much of the foundations of electromagnetism were discovered through work by Oersted, Ampere, Faraday, Gauss, and Maxwell. After all of the work in developing classical physics, electromagnetism, and statistical mechanics one would expect that a theory of magnetism would be possible, however this is not the case. In his PhD thesis Bohr (and independently by van Leeuwen) showed that the logical conclusions of these theories is that the time-averaged magnetisation of any system in thermal equilibrium is always zero. The so-called Bohr-van Leeuwen theorem contradicts our everyday experience with such rigour that we are forced to consider that we do not fully understand the nature of the problem.

Coincidentally, around the same time as the Bohr-van Leeuwen theorem was formulated the foundations of quantum mechanics were being formed. With the knowledge of the electron's intrinsic spin we are able to formulate a theory of magnetism which is consistent with our experience. One could be forgiven for not anticipating that something so macroscopic and prevalent as magnetism to be rooted in something so unfamiliar and strange as the spin of an elementary particle. However in some sense we have replaced one mystery with another as the fundamentals of what spin *is* is not well understood.

Continuing down the path of investigating the nature of magnetism in ever increasing detail, we wish to study magnetic dipole-dipole interactions between atoms at ultralow temperature in order to study spin interactions and magnetism within a quantum system. We have opted to realise this system in a dilute, weakly interacting gas near absolute zero temperature, forming a state of matter called a Bose-Einstein Condensate (BEC). The behaviour of BECs is well described in most cases by now due to extensive theoretical and experimental study since their concep-

tion in 1925 and first realisation in 1995. One regime which remains to be explored experimentally is how BECs behave in ultralow magnetic fields.

BECs can be prepared entirely optically, leaving the spin of the atoms as a degree of freedom along with their translational degrees of freedom leading to a rich dynamical evolution of the quantum state[6–12] Such a system realises a mathematical spinor and is thus termed a spinor BEC. The ambient magnetic field strength in all previous BEC experiments has been so large that degeneracy of the hyperfine Zeeman sublevels is lifted to such a degree that there is no coupling between them. Lowering the magnetic field such that the interaction energy due to dipole-dipole interactions within the system are comparable to the sublevel splitting allows these interactions to drive population transfer between these states and in turn affect the dynamical evolution of the spinor BEC. These dipole-dipole interactions give rise to a further layer of richness and complexity to the dynamical behaviour of the spinor BEC. We will refer to a spinor BEC in this regime as a dipolar BEC.

In all previous experiments the energy difference between states with different spin projections, called the Zeeman splitting, has been too large to allow dipole-dipole interactions between atoms to induce population transfer between states. This energy difference can be lowered by reducing the ambient magnetic field. By lowering the magnetic field to sufficiently low levels it becomes possible to probe quantum magnetic phenomena within a BEC by studying state populations (via optical density) and spatial distributions. The field levels required to reach this regime can be estimated based on the ~1 nK dipole-dipole interaction energy per atom and the 0.7 MHz/G ground state Zeeman splitting for ^{87}Rb[13] of ~300 pT (= 3 μG). This requirement is critical to the success of our experiment and as such we have gone to great lengths to achieve ultralow magnetic fields in our setup, with particular care taken in the region around the BEC.

To realise a sufficiently weak magnetic field around the BEC to achieve a dipolar BEC we require to shield the ambient magnetic field to a very high degree and

stabilise it to prevent transient fields causing the system to quench. This can be done using passive magnetic shielding, active magnetic shielding, or a combination of both. Reducing the ambient magnetic field to a level where these magnetic interactions affect the system allows us to study persistent mass- and spin-current ground states[14–16], the Einstein-de Haas effect in a quantum system[17–20], geometrically frustrated spin systems (such as spin ice on a square lattice[21–24] as well as in other geometries[25–28]), prepare macroscopic entangled states (in a single trap[29] or on a lattice[30]), search for a quantum-to-classical transition in spin space[31], perform state-of-the-art magnetometry[32,33], and potentially even study the transport and properties of effective magnetic monopole excitations in a well-controlled manner[15,26,27,34–36].

This thesis presents the basic theory behind atomic physics experiments (chapter 2), the design and construction of our vacuum chamber (chapter 3, section 3.1), the optical considerations of our experiment (chapter 3, section 3.2), the theory behind and design of a passive magnetic shield (chapter 3, section 3.3), and finally the design and construction of an active magnetic shielding system (chapter 4). The key results of this thesis are the design of a passive magnetic shield which can reduce the ambient magnetic field by a factor of $> 10^8$ in 3D and the implementation of a high performance active magnetic shielding system for atomic physics experiments, which has been replicated by collaborators at the Institut für Laserphysik of Universität Hamburg and is also intended to be installed in collaborators' labs at the Laboratoire Kastler Brossel of École Normale Supérieure. The design of our vacuum chamber, UHV glass cell, 2D MOT telescopes and peripheral hardware, evanescent fibre-coupler network, passive magnetic shield, and active magnetic shielding system was undertaken solely by me. The initial concepts and implementation of the active shielding system were undertaken as part of a secondment in the group of Dr. Fabrice Gerbier at the Laboratoire Kastler Brossel of École Normale Supérieure under the direction of Dr. Tilman Zibold. A publication of the design, implementation,

and performance of the active magnetic shielding system described in this thesis and its integration into an ultracold atoms experiment is currently being prepared with collaborators at the Institut für Laserphysik of Universität Hamburg and the Laboratoire Kastler Brossel of École Normale Supérieure.

2

Theory overview

This chapter aims to present the basic atomic physics requisite to understand the core of this thesis. Additional topic-specific theory will be presented alongside the topic itself so that individual chapters of this thesis are self-contained.

Atomic configuration 2.1

A given element comprises a nucleus composed of protons and neutrons which is surrounded by electrons. The electrons arrange themselves into energetically favourable "shells", obeying Hund's rule, where "complete" shells (i.e. those in which every available quantum state is occupied) are favoured. Electrons left over from this filling of closed shells form the outermost shell, which is incomplete except for the

case of the noble gases. The electrons in this shell are responsible for the reactivity of an element and to a large extent determine its spectroscopic properties. Outer shell electrons have well-defined energy levels to which they can be excited via absorption of photons. In the case of the alkali metals the outer shell contains only a single electron which greatly simplifies the energy level structure of the atom. This experiment, like many others, uses an isotope of rubidium (^{87}Rb) due to the sophisticated supporting technology available and the favourable properties of Rb for laser cooling. Some of the first cold atoms experiments were performed on Rb due to its main spectroscopic lines being in a region of the near infrared spectrum where low cost, good quality laser diodes were available due to being mass-produced for use in CD and DVD players.

As reaching ultralow magnetic fields is central to our experiment, our atoms will almost exclusively be in the low field limit, where J and F are good quantum numbers. Notation to describe the configuration of electrons within the atom can be simplified greatly in the case of alkali metals by omitting all but the outer electron. Energy levels of the atom can then be labelled with the spectroscopic notation which defines the state of the outer electron, given in the form

$$n^{(2S+1)}L_J.$$

This compact notation completely describes the quantum state of an electron via the quantum numbers n, S, L and J, where n is the principal quantum number (detailing the shell in which the electron is situated), S is the total electron spin angular momentum ($S = |\mathbf{S}| = \frac{1}{2}$ for alkalis as the filled inner shells have $S = 0$), $L = |\mathbf{L}|$ is the Azimuthal quantum number which specifies the orbital angular momentum of the state (with the letters $s, p, d, f \ldots$, continuing alphabetically, corresponding to the values 0, 1, 2, 3...) and $J = |\mathbf{J}| = |\mathbf{L} + \mathbf{S}|$ is the total electronic angular momentum. These quantities describe the *electronic* properties but to fully describe

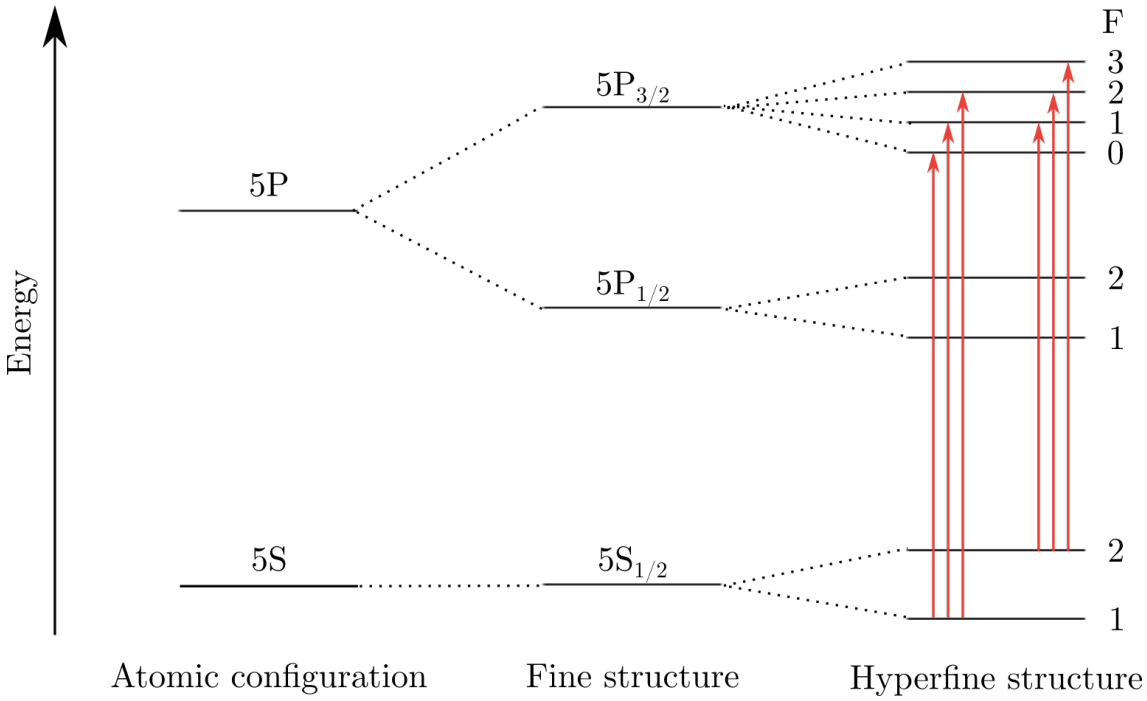

Figure 2.1 The energy level structure for ^{87}Rb showing allowed transitions (*red*). On the left we have the most simplistic view of the energy levels and progressing to the right we include more and more contributions to the energy levels, ultimately giving us a fairly accurate picture of the atomic energy structure. Note that the spacings are exaggerated for clarity, as the hyperfine splitting (ca. 500 MHz for the excited state and ca. 6.8 GHz for the ground state, both $\ll 1$ *nm* difference in wavelength) is much smaller than the fine splitting (ca. 15 nm difference in wavelength).

the energy levels we also need to take into account the *nuclear* angular momentum, $I = |\mathbf{I}|$ (for ^{87}Rb, $I = \frac{3}{2}$). Inclusion of the nuclear spin gives the total angular momentum of the atom, $F = |\mathbf{F}| = |\mathbf{J} + \mathbf{I}|$.

The ground state of the outer electron in ^{87}Rb is the $5^2S_{1/2}$ state. The relative orientations of the electronic and nuclear spins cause this state to split into two levels with $F = 1$ and $F = 2$, with the anti-aligned $F = 1$ state having lower energy. Applying the same reasoning to the first excited state of the electron one obtains the level structure illustrated in figure 2.1.

Not all transitions between these levels are allowed, but only those that obey the

selection rules

$$\Delta J = 0, \pm 1, \tag{2.1}$$

$$\Delta L = 0, \pm 1, \tag{2.2}$$

$$\Delta S = 0. \tag{2.3}$$

A transition must obey all of these rules simultaneously. Equations 2.1 and 2.2 together relate to the conservation of momentum when a photon is absorbed by the atom. The photon's momentum is transferred to the atom, which requires to change its arrangement of $\mathbf{L}, \mathbf{S}$ and $\mathbf{I}$ vectors to account for this change in momentum. As one typically only considers the total angular momentum of the atom in practice it is more convenient to use a more compact version of the selection rules, $\Delta F = 0, \pm 1$.

Atomic transition lines are typically characterised by their "natural" linewidth, Γ_0, which is the full-width at half-maximum (FWHM) of their emission spectrum at zero temperature ($\Gamma_0 \approx 6$ MHz for the D_2-line of ^{87}Rb). A natural linewidth is defined because the observed line can be broadened by temperature, pressure, laser intensity and other factors.

2.2 The magneto-optical trap

In order to study the properties of a cloud of atoms we must cool them so that they move very slowly and hold them in a defined region to prevent them from drifting away. Both of these can be achieved using the ubiquitous Magneto-Optical Trap (MOT) which has become almost a requirement for experiments on cold atoms. The phenomena which this trap relies on are presented in order below.

The Doppler shift 2.2.1

Crucial to all cold atoms experiments, the Doppler shift is an apparent change in frequency, Δf, of radiation from a source, with natural frequency f_0 and propagation velocity v_{rad}, due to relative motion between the observer, travelling at v_o, and the source, travelling at v_s, which is given by

$$\Delta f = \left(\frac{v_o - v_s}{v_{rad}} \right) f_0. \tag{2.4}$$

Here the direction of positive velocity is defined as pointing from the observer towards towards the source. As this effect is scaled by the velocity of the radiation it can be readily encountered with audio signals, where $v_{rad} \approx 360$ ms^{-1} is comparable to common source/observer velocities, such as when an ambulance siren appears to decrease in pitch when driving past a stationary observer. Such effects also occur with light and can be exploited to exert a powerful cooling effect on atomic gases. Whilst the relative frequency shifts for light are around 10^6 smaller than those of audio signals the absolute frequency shifts are larger due to the higher fundamental frequency of light.

To get a feel for how much this shift is for room temperature ^{87}Rb atoms we can use kinetic theory to determine the root-mean-squared (RMS) velocity of atoms, $\overline{v}$, at a given temperature, T, using

$$\overline{v} = \sqrt{\frac{3kT}{m}}, \tag{2.5}$$

where k is Boltzmann's constant and m is the mass of the particle (for ^{87}Rb, $m \approx 1.443 \times 10^{-25}$ kg[13]). At $T = 300$ K, equation 2.5 gives $\overline{v} \approx 290$ ms^{-1} and leads to a Doppler shift of $\Delta f \approx \pm 370$ MHz for 780 nm light (resonant with the $5S \rightarrow 5P$ transition) where the sign depends on the relative direction of motion of the atom. In a thermal vapour there will be a Gaussian velocity distribution of the atoms

around $\overline{v}$ and this will cause "Doppler broadening" of the transition lines. This has consequences for traditional spectroscopic techniques for example, as absorption lines will be Doppler broadened from $\Gamma_0 \approx 6$ MHz to $\approx 2\Delta f = 740$ MHz obscuring any detail of the hyperfine structure (which spans a total of $\approx$ 500 MHz in the ^{87}Rb $5P_{3/2}$ manifold[13]). A further concern is that this shift is $> 60\ \Gamma_0$ and will cause a moving atom to be far off resonance with a stationary laser tuned to an atomic transition frequency. However, one can counteract this effect by tuning the stationary laser's frequency to be less than that of the atomic transition ("red-detuned"[1]) such that the Doppler shift for an atom travelling towards the laser brings the observed frequency closer to resonance in the atom's frame of reference, whilst conversely atoms moving away are shifted even further from resonance. This gives rise to a direction-dependent absorption of light for moving atoms and forms the basis for cooling atomic gases with lasers.

Optical molasses 2.2.2

Using the direction-dependent absorption of red-detuned light due to the Doppler shift, one can engineer a method of cooling an atomic gas using several retroreflected beams. To illustrate this we can take a 1D example of two laser beams which are counter-propagating through an atomic gas and are red-detuned from some atomic resonance, as shown in figure 2.2. As discussed in section 2.2.1, atoms travelling towards either laser beam preferentially absorb photons from that beam. Conservation of momentum requires that the atom receives a small increase in momentum (a "momentum kick") from the absorption of each photon in the initial direction of travel of the photon (in this case, away from the laser). This means that as an atom travels towards the laser it absorbs photons which subsequently reduce its velocity in that direction, effectively reducing its temperature.

[1] Off-resonant light is said to be *detuned* from resonance and is classed as either red- or blue-detuned for below or above resonant light. One can imagine that a green light source would appear more red if its frequency were reduced and more blue if it were increased.

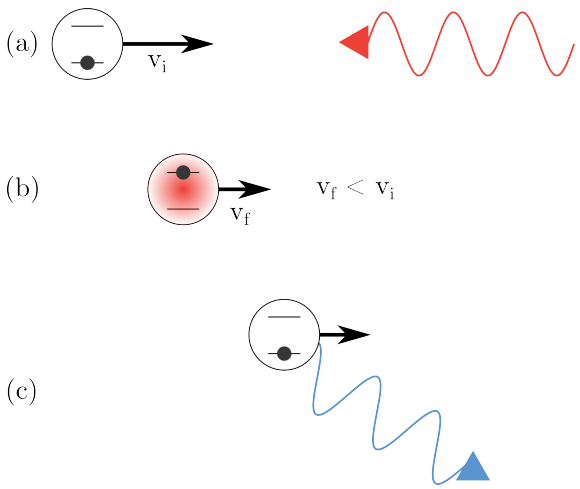

Figure 2.2 A visual explanation of laser cooling. (a) Initially, an atom and red-detuned photon from a laser beam are travelling towards one another such that the atom sees a Doppler-shifted photon near resonance. (b) The Doppler-shifted photon is absorbed, placing the atom in an excited state and reducing the atom's velocity due to conservation of momentum. (c) After a short time a photon is spontaneously re-emitted by the atom bringing it back to the ground state, but as the emission direction is random the time-averaged momentum kick is zero and the atom experiences a net force along the axis of the laser beam.

Of course, soon after a photon has been absorbed by an atom it is re-emitted via spontaneous emission giving the atom another momentum kick of equal magnitude to the absorption of a photon. However, as the direction of the spontaneously emitted photon is random the momentum kick averages to zero over many cycles whereas the absorbed photon kick, in a constant direction which reduces the atom's velocity, remains and leads to a net cooling effect on an atom in the axis of the beam. Extending this idea to 3D by using 3 mutually orthogonal counter-propagating beam pairs which intersect to produce a volume in which the average velocity of an atom is reduced in all 3 spatial directions. In this overlap volume the atoms experience a velocity-dependent force acting upon them, much like that of a particle travelling in a viscous medium, causing researchers to name this overlap volume an "optical molasses".

The velocity reduction of an atom due to absorption of a single photon in 1D can be calculated by considering the momentum of the atom before, p_i, and after,

p_f, the absorption process

$$\begin{aligned} p_i &= p_{atom} = mv_i, \\ p_f &= p_{atom} - p_{photon} = mv_i - \frac{h}{\lambda} = mv_f, \\ \therefore \Delta v &= v_i - v_f = \frac{h}{m\lambda}, \end{aligned}$$

where m is the mass of the atom, h is Planck's constant and λ is the wavelength of the photon. For ^{87}Rb this gives the result that each photon absorption reduces an atom's velocity by around 6 mm s^{-1}. This is only a small fraction of a ^{87}Rb atom's average velocity at room temperature of ~290 ms^{-1} however this process can occur several million times per second allowing for significant and rapid cooling of atoms using this technique. It should be noted that this does not *trap* the atoms but merely slow them down within a certain volume. In order to trap atoms we require that a position-dependent (rather than velocity-dependent) force is exerted upon them. An optical molasses can be turned into a trap for atoms with the addition of a magnetic field gradient or a considered choice of the laser beam parameters which will be discussed in the following sections.

The Zeeman effect 2.2.3

In a magnetic field each of the hyperfine energy levels is split further into $2F + 1$ components pertaining to the relative orientations of the total angular momentum and the external magnetic field, $\mathbf{B}$. The energy splitting, ΔE, associated with this is given by[37]

$$\Delta E = \mu_B g_F m_F B \tag{2.6}$$

where μ_B is the Bohr magneton, g_F is the Landé g-factor for a given hyperfine level and m_F is the projection of $\mathbf{F}$ onto the axis of the magnetic field $\mathbf{B}$ (and following convention, $B = |\mathbf{B}|$). This modification of the energy levels of different m_F states,

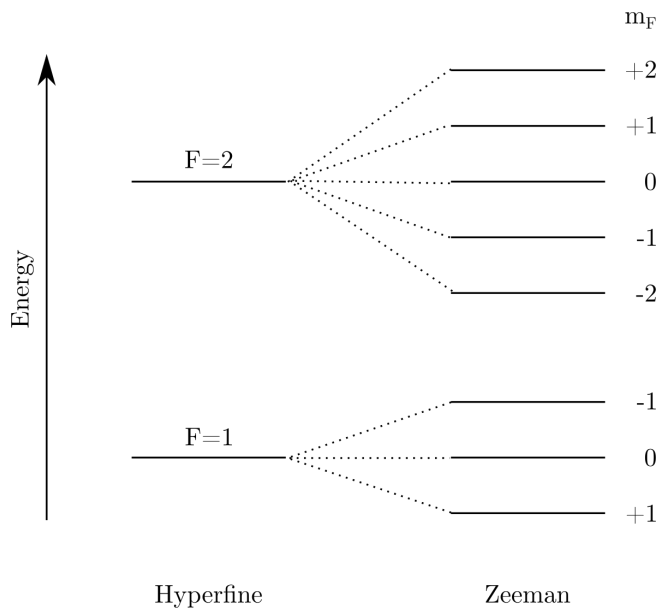

Figure 2.3 ^{87}Rb hyperfine levels split by an energy ΔE in a magnetic field $\mathbf{B}$ due to the Zeeman effect. Note that the order of the Zeeman sublevels (from lowest to highest energy) are inverted between the two levels due to the difference in sign of the g-factor between the two levels.

shown in figure 2.3, is known as the Zeeman effect and can be used to trap atoms.

As the Zeeman shift of an m_F state depends on the local value of $\mathbf{B}$, by applying a spatially-varying magnetic field one can engineer a spatially-dependent energy level shift for that state. For example, by generating a quadrupole field (e.g. $\mathbf{B}(x, y, z) = \alpha x\hat{\mathbf{x}} + \alpha y\hat{\mathbf{y}} - 2\alpha z\hat{\mathbf{z}}$) one can create a situation where the Zeeman shift is zero only at the origin with linearly increasing shifts outward from this point. Akin to the operating principle of the optical molasses discussed above one can send counter-propagating red-detuned laser beams through the zero field region and observe selective absorption of photons from one beam over another. Here, the selectivity arises due to the atoms shifting toward or away from resonance by the spatially-dependent Zeeman shift they experience (rather than by their Doppler shift, as in an optical molasses). For states with $m_F < 0$ in the $F = 1$ maniford the resulting force applied to the atoms through photon absorption points towards the zero field point from every position with $\mathbf{B} > 0$ (conversely for $m_F > 0$ states where

$\mathbf{B} < 0$), providing a restoring force to the zero field region where atoms are trapped.

When describing a MOT on paper it is convenient to use the notation of σ^+ and σ^- in order to describe the circular polarisations due to their implicit angular momentum of $L = \pm 1$. However, this notation is not such a useful one to use in practice (i.e. in the lab) as it relies on having a known magnetic field vector to set the quantisation axis of the angular momentum. It should be clarified here that at all times we use the convention of defining the polarisation with a magnetic field parallel to the Poynting vector. For example, an electric field vector which rotates clockwise when looking along the axis of propagation would be termed σ^+ or *right-handed*, similarly σ^- is *left-handed* (following the same convention as screw threads). Specific m_F states can be addressed through the increased absorption probabilities for certain states based on the polarisation of the incident photons. This description is summarised in figure 2.4, which shows the spatial energy level splitting and correct beam polarisations to address m_F states which can be trapped. The combination of spatial magnetic field gradient with an optical molasses (having correctly chosen polarisations) overlapped at the point of zero field is called a Magneto-Optical Trap (MOT) and is used in almost all experiments on cold atoms. The optical molasses provides a slowing of thermal atoms whereupon they can be trapped by the much weaker magnetic trap.

Due to Maxwell's 3$^{\text{rd}}$ law ($\nabla \cdot \mathbf{B} = 0$) we have a defined "strong" axis and two "weak" axes in a MOT. This can be understood by considering a quadrupole field gradient applied along the z-axis. In order to keep $\frac{\partial \mathbf{B_x}}{\partial x} + \frac{\partial \mathbf{B_y}}{\partial y} + \frac{\partial \mathbf{B_z}}{\partial z} = 0$ one field configuration would be $\frac{\partial \mathbf{B_x}}{\partial x} = \frac{\partial \mathbf{B_y}}{\partial y} = -\frac{1}{2}\left(\frac{\partial \mathbf{B_z}}{\partial z}\right)$, thus the field gradient is naturally lower in two axes and more importantly changes sign. This gives rise to the fact that the polarisation choices for the strong axis are different to those for the weak axis — i.e. if σ^+ light travlleing along +ve x and y provide cooling then we should choose σ^- light to travel along +ve z in order to provide cooling in all axes of the MOT. We should note here that this works only for a right-handed co-ordinate system.

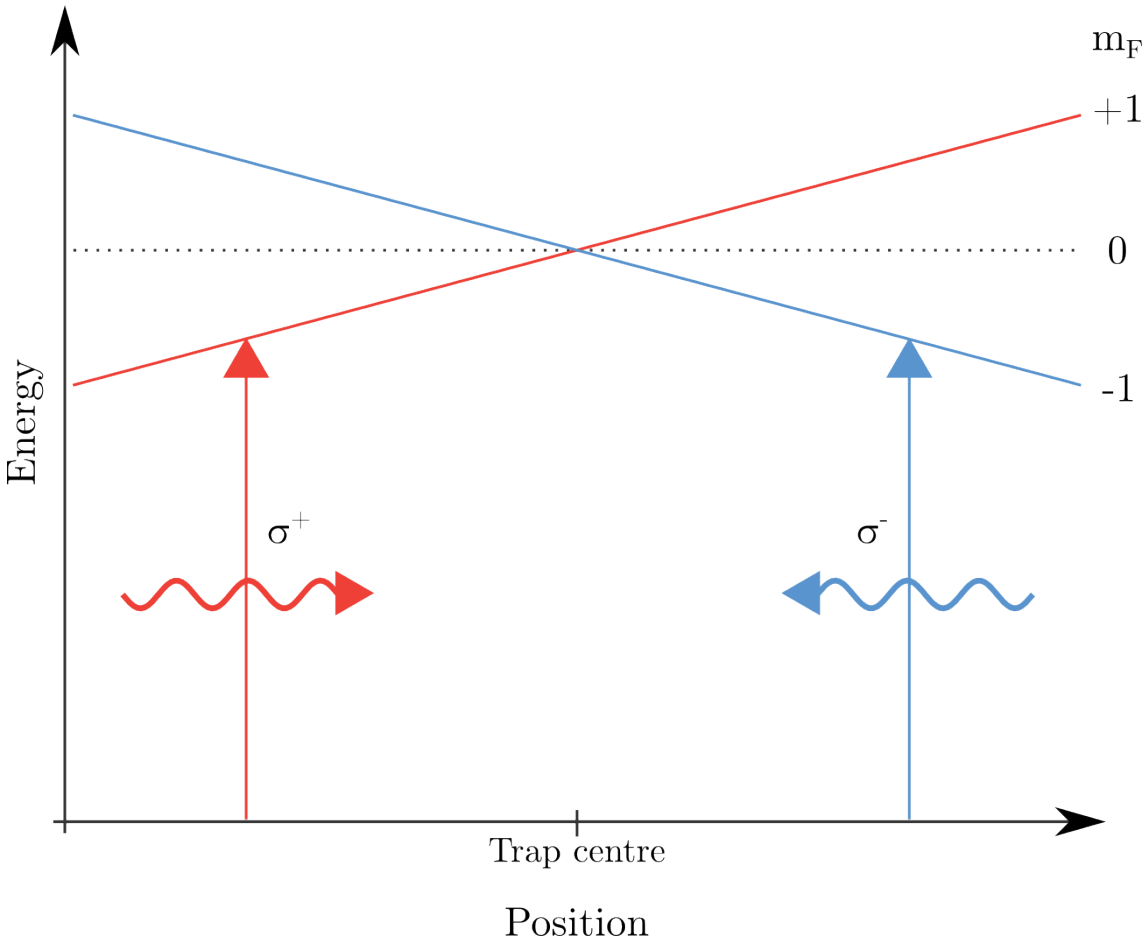

Figure 2.4 The energy levels of a cooling transition which are split linearly in space due to the Zeeman effect in a quadrupole magnetic field, which is linear in space from the centre of the trap. Light coming from the right of the graph (with a polarisation of σ^-) will interact with the atoms in the $m_F = -1$ levels, not those in $m_F = +1$ or $m_F = 0$ and vice-versa for the left hand side of the graph. This causes the atoms to preferentially absorb one photon or the other depending on which direction they are travelling which pushes them towards the centre of the trap, providing a position dependent force. NB: this graph is drawn with a single quantisation axis in mind, whereas in reality the sign of the field changes on either side of the trap centre, which would cause the m_F levels to change sign on one side of the trap centre.

Using a circular analyser (comprising a $\lambda/4$ waveplate followed by a polarising beamsplitter at 45°) we can deduce that a beam is *either* one handedness of polarisation or the other by whether the analyser allows the light to pass straight through or is reflected out of the orthogonal port of the beamsplitter. Whilst we can separate the handednesses of the beams we have no easy way of determining *which* handedness we are dealing with. In addition, along each axis the magnetic field vector changes sign which flips the convention for σ^+ and σ^- beams depending on whether the beam is travelling parallel or anti-parallel with the axis. E.g. if we have a +ve field gradient along x (i.e. $\mathbf{B} \propto x$) then the beam travelling along +ve x is aligned with the local magnetic field vector and we can use our analyser to select one handedness for this beam. However, the magnetic field is anti-parallel with the beam travelling along −ve x and thus the convention for determining $\sigma^\pm$ is inverted. Thus in order to make this beam have the opposite σ polarisation to the beam travelling along the +ve x direction we must use our analyser to select *the same* handedness of light as for +ve x. Here is where the convention of σ^+ and σ^- breaks down in practice. This means that the beams comprising the "weak" axes should all have the same handedness and conversely the beams comprising the "strong" axis should have opposite handedness to this in order to produce the correct polarisations for laser cooling.

The dipole trap 2.3

In order to further cool and increase the density of our atomic cloud beyond that achieved by our MOT we require a different trapping mechanism to that of the MOT. Magnetic traps, essentially a much higher field version of a MOT, are commonly used in cold atoms experiments to evaporatively cool to BEC. We however wish to avoid the use of a magnetic trap as this would likely cause our passive magnetic shield to become magnetised, which greatly affects its performance. Trapping atoms using

the dipole force from an intense laser beam and how this can be used to cool atoms is detailed in this section.

The dipole force 2.3.1

The electric field from a sufficiently high power laser beam is able to polarise atoms, giving them an effective dipole moment which oscillates at the frequency of the incident laser field. This induced dipole moment interacts with the non-uniform potential created by the electric field of the laser beam due its spatial profile. For a red-detuned laser beam the minimum of this potential is coincident with the peak intensity, meaning that atoms in such a beam become polarised and are then pushed towards the highest intensity region of the beam (conversely, they are pushed away from the highest intensity for a blue-detuned beam). A common analogy used to describe this behaviour is that of a small dielectric sphere placed in a laser beam with a Gaussian intensity profile. The sphere acts like a concave lens with a short focal length and experiences a net force due to the unbalanced forces of light refracted from the centre of the beam (more intense) and the edges of the beam (less intense), shown in figure 2.5, which causes the sphere to be drawn to the centre of the beam. By focusing the beam we can extend this trapping from 2D to 3D. To avoid interacting resonantly with the atoms, whereupon we scatter photons which causes heating of the cloud, the detuning of these beams must be very large. In our case we use 1550 nm beams, which yield a detuning of around $3 \times 10^7\ \Gamma_0$ from the D_2 line. Such a large detuning is chosen due to the availability of low cost, compact, high power lasers at 1550 nm due to their widespread use in the telecoms industry. As the detuning of these beams is so large compared to typical thermal drifts of the laser frequency we can neglect these and thus we do not require this laser to be locked to an atomic transition.

The time-averaged force, $\overline{\mathbf{F}}$, exerted on an atom in a laser beam can be shown

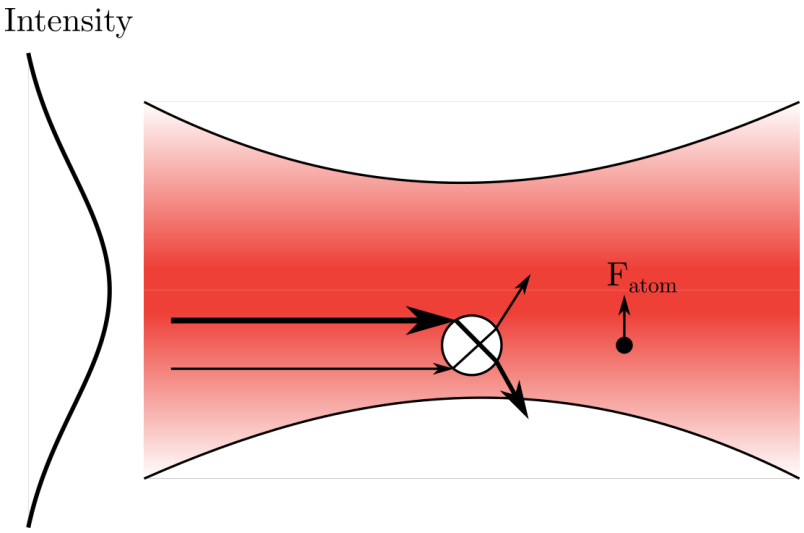

Figure 2.5 A schematic diagram of how the dipole force attracts atoms to the centre of an intense, far red-detuned laser beam. The unbalanced force from the refraction of the intense central region of the beam compared to the refraction of the much less intense wings of the beam give rise to a restoring force towards the centre of the beam. In this diagram, we only show the restoring force radially along the beam, however this beam is also focused (so that the intensity also varies with a Gaussian profile *along* the beam) and this introduces another trapping field axially ensuring the atoms are confined in 3D.

to be[37]

$$\overline{\mathbf{F}} = \frac{e^2}{2m\epsilon_0 c}\frac{1}{\delta^2 + \left(\frac{\beta}{2}\right)^2}\left\{-\delta\frac{\nabla I}{2\omega} + \left(\frac{\beta}{2}\right)\frac{I}{\omega}\hat{\mathbf{k}}\right\}, \tag{2.7}$$

where e is the electron charge, ω is the angular laser frequency, δ is the detuning from an atomic resonance, β is the so-called "damping ratio" of the trap, I is the intensity of the beam and k is the wavevector. The damping ratio is defined as $\beta = -\hbar k^2 \frac{4s_0(\delta/\gamma)}{[1+s_0+(2\delta/\gamma)^2]^2}$ for a travelling wave[38], where $s_0 = I/I_s$, where I_s is the saturation intensity for a given transition. The first term in equation 2.7 is called the dipole force and the second term is called the scattering force. The scattering force is used to trap atoms in a MOT and the dipole force enables trapping in the focus of a high power laser beam. Ideally the scattering force should be low in order that the dipole force dominates this equation, hence we want a small value of β. This can be achieved with large δ and I, as in our experiment where a 10 W 1550 nm laser is used for our dipole trap.

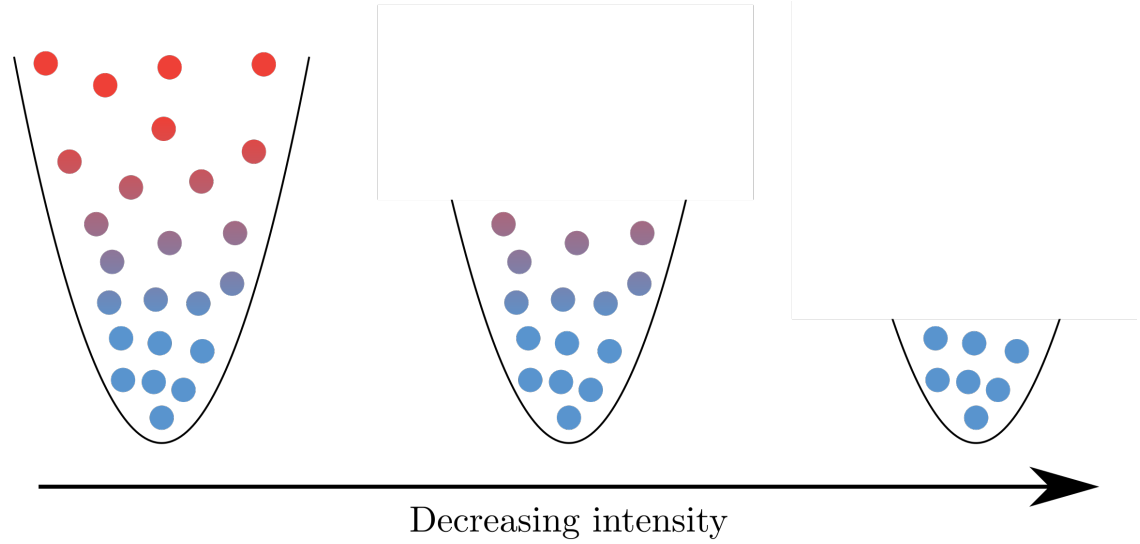

Figure 2.6 A schematic representation of evaporative cooling in a dipole trap. The initially deep trap (*left*) created by a high intensity laser beam contains a wide range of different temperature atoms. As the intensity is reduced, the trap depth is lowered, allowing the hottest atoms to escape (*centre*). Finally, after several stages of this process we achieve a cold, dense cloud in a relatively shallow trap (*right*). Note that the intensities given here are only indicative.

2.3.2 Evaporative cooling

The dipole trap itself does not cool the atoms as the trap potential is conservative. In order to cool atoms in a dipole trap we remove the hottest atoms from the trap, allow the remaining atoms to re-thermalise (to a now lower temperature) and repeat this process many times. This process is called evaporative cooling and can be achieved by initially using a high power laser beam, which creates a deep potential well for the atoms, then ramping down the beam power in stages. At each stage the hottest atoms from the previously deeper trap have enough kinetic energy to escape the now shallower trap, demonstrated visually in figure 2.6. This new shallow trap is maintained for some time in order to allow the atoms to re-thermalise and the process repeats. However, this process is *extremely* lossy, ejecting around 99.9% of our trapped atoms in order to only retain the coldest fraction.

Optical lattices 2.4

An optical lattice in 1D is a standing wave generated by two counter-propagating laser beams. This creates a periodic series of potential minima in electric field spaced at $\lambda/2$ where atoms can be trapped. Like a dipole trap, this has to be operated far detuned from resonance in order to not resonantly interact with the atoms. Extending this idea to 2D we can create an optical "egg box" where we can trap atoms on a grid, and in 3D we can trap atoms at the sites of an effective cubic crystal lattice.

Whilst optical lattices can be used as a cooling technique[39], in our setup we wish to use this to enhance and ultimately tune the dipolar interaction strength in our BEC. We can achieve this via changing the so-called shape anisotropy energy via altering the shape of the traps at each site. When each axis of the trap has equal optical power the trap is said to be isotropic, producing spherical isopotentials for atoms at each lattice site. Unequal optical powers in each axis deforms the isopotential at each lattice site towards either oblate (pancake-shaped) or prolate (cigar-shaped) spheroids, which provides anisotropic confinement. This modification of the trapping anisotropy alters the ground state of atoms at each lattice site due to intrasite dipole-dipole interactions. For example, in an oblate trap the dipole-dipole interaction energy (E_{dd}) between atoms is minimised when their dipoles are confined to the plane of the trap (i.e. the surface of the pancake) whereas in a prolate trap the E_{dd} is minimised when atoms align along the axis of the trap (i.e. the axis of the cigar). This mechanism causes the effective dipolar interaction strength to be greatly enhanced in an oblate trap[14,16,40]. However, with an array of prolate traps we can observe an almost purely dipolar system over the lattice due to the fact that at each site the atomic dipoles tend to align creating a large collective spin at each site. As the dipole-dipole interaction has a much longer range than other types (which are mostly collisional) these collective spins are able to interact with those at

other lattice sites, even beyond their nearest-neighbour sites, giving rise to a lattice of collective spins whose main interaction is dipolar[41,42].

2.5 Bose-Einstein condensation

Once an atomic cloud of bosons becomes cold and dense enough a quantum phase transition occurs forming a Bose-Einstein Condensate (BEC). A BEC is essentially an entire cloud of atoms in the quantum ground state of the system. Within the limts of the Thomas-Fermi approximation, one signature of BEC formation which demonstrates this is that the spatial distribution of the atoms changes to mirror the trapping potential (usually parabolic) rather than a thermal Gaussian distribution. Conveniently, a BEC is sufficiently large that it is macroscopic (~100 μm in length — around the thickness of high grade paper) and thus can be viewed on a video camera with only modest magnification.

Whilst a full derivation of the BEC transition is rather lengthy, one can describe the basic idea of a BEC using the so-called de Broglie wavelength. As an extension of the paradigm of wave-particle duality, de Broglie proposed that all matter has a characteristic wavelength which is proportional to its mass, m, and temperature, T, known as the de Broglie wavelength,

$$\lambda_{dB} = \frac{h}{\sqrt{2\pi m k_B T}}. \tag{2.8}$$

What is important to note about equation 2.8 is that this wavelength increases with decreasing temperature. When the de Broglie wavelength becomes comparable to the spacing between the atoms in the cloud the notion of an individual particle becomes poorly defined and wave-like effects begin to become manifest in the system. Once the density, ρ, and de Broglie wavelength reach the criterion that

$$\rho = \frac{2.6}{\lambda_{dB}^3}, \tag{2.9}$$

we see the transition into a BEC[37]. As a more intuitive picture, re-arranging equation 2.9 we can see that this criterion can be written as

$$\therefore N\lambda_{dB}^3 = 2.6\ V,$$

where N and V are the number of atoms and V is the volume of the cloud respectively. The left hand side of equation 2.5 can be thought of as the "quantum volume" of the atoms based on their de Broglie wavelength and states that this "quantum volume" needs to be equal to (or greater than) 2.6 times the physical volume of the atoms in order to create a BEC. This leads to the de Broglie wavelengths of each atom overlapping with each other almost entirely, causing the whole cloud to behave as if it were a single quantum object described by a single wavefunction.

Interactions within a condensate 2.5.1

Within our BEC there will be many interactions between the atoms due to collisions, interactions with external magnetic fields, and interactions between atomic dipoles. These interactions all have different energy scales, and it is the latter interactions that we are most interested in. Expressed in units of temperature ($E_{(K)} = E_{(J)}/k_B$), collisional interactions dominate with an energy of ~100 nK whilst the dipolar interaction between atoms is only ~0.1 nK[15]. One would intuitively expect, as was assumed for a long time, that we should not see any effects from the dipolar interactions as they are such a small part of the total interaction energy, however as the atoms (being bosons) act collectively below the BEC transition these dipolar interactions become enchanced, allowing them to be more readily measured[15].

Collisional interactions can either be spin-dependent or spin-independent based on whether or not they affect the internal states of the atoms. Spin-independent collisions leave the spin projection of the atoms involved in the collision unchanged. These interactions provide the largest contribution to the total collisional energy

with an energy of around ~100 nK[15]. Spin-dependent collisions can change the spin projection of the atoms involved in the collision, but the overall m_F is conserved (i.e. if two $m_S = 0$ atoms collide they can emerge with a net spin of zero in any configuration, such as $m_{S1} = +1, m_{S2} = -1$). This allows for population transfer between states without changing the overall magnetisation of the BEC. Such interactions contribute ~1 nK to the interaction energy[15]. Finally, dipolar interactions can also change the spin projection of the atoms involved in a collision, but these collisions *do not* conserve the total m_F. Such interactions lead to spontaneous magnetisation or demagnetisation of a BEC[43,44] and can be used in combination with optical pumping as a cooling technique[45]. These interactions contribute a mere ~0.1 nK to the interaction energy[15]. The discussed interactions are summarised visually in figure 2.7.

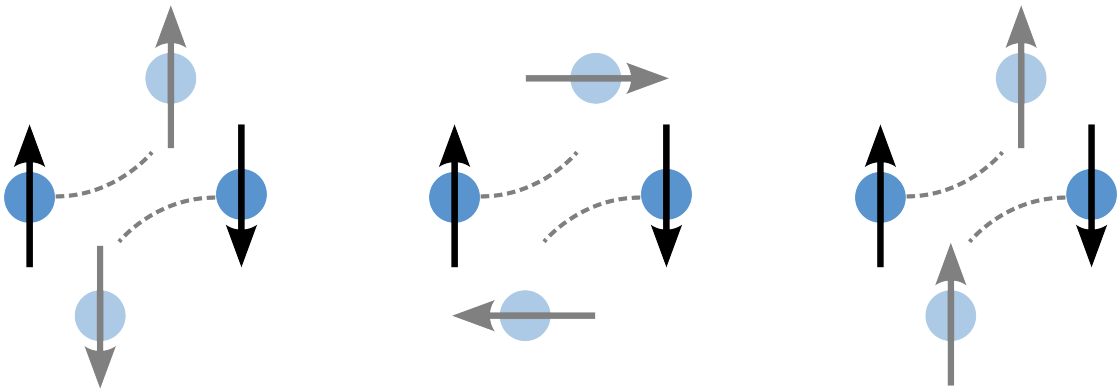

Figure 2.7 Representations of the types of interactions we will see in our condensate: (a) spin-independent collisions which don't alter the orientation of the atomic spins, (b) spin-dependent collisions which can alter the orientation of the atomic spins, but conserve the overall m_F and (c) dipolar interactions which can alter the orientation of the atomic spins *without* conserving the overall m_F.

However, these energies are miniscule in comparison with interactions with an external magnetic field. Using equation 2.6, and the fact that $g_F = -\frac{1}{2}$ for the $5S_{1/2}(F = 1)$ state of ^{87}Rb[13] (i.e. the ground state) gives us a Zeeman interaction energy of $-16.8\ \mu$K/G. As interactions with an external field are on the order of ~μK

and collisional interactions are on the order of ~nK we can see that this external fields govern the behaviour of our BEC. As we wish to study the dipolar interactions within our BEC, we require to reduce these interactions with an external field to a level lower than that of the dipolar interactions — which means reducing the ambient magnetic field to ≤ 1 μG. This requirement is critical to the success of our experiment and as such we have gone to great lengths to achieve ultralow magnetic fields in our setup, with particular care taken in the region around the BEC.

Spinor BEC 2.5.2

BECs are often produced in state-selective magnetic traps which cause the BEC to be composed of only a single spin projection. Once the BEC threshold has been crossed the internal spin degree of freedom is effectively "frozen" meaning that spin-dependent collisions are precluded. However, BECs can also be formed in state-agnostic optical traps which can produce a BEC composed of mixture of spin states. Such a BEC can be described as n overlapping "pure state" BECs using a vector order parameter in the wavefunction to represent the population in each spin state. This type of BEC is called a *spinor* BEC due to the fact that it represents a mathematical spinor. Having the atoms' spin projection as a degree of freedom enables the extra mode of spin-dependent collisions, making the dynamical properties of spinors more rich and complex than those of a regular, spin-frozen BEC.

Dipolar BEC 2.5.3

Whilst a spinor BEC has the atomic spin projection as a degree of freedom it typically does not exhibit magnetic phenomena. In all previous experiments the energy gap between states with different spin projections has been too large to allow dipole-dipole interactions between atoms to induce population transfer between states. The energy gap between these states is proportional to the local magnetic field strength

and thus can be lowered by reducing the magnetic field. By lowering the magnetic field to sufficiently low levels it becomes possible to probe quantum magnetic phenomena within a BEC by studying state populations (via optical density) and spatial distributions. The field levels required to reach this regime can be estimated based on the ~1 nK dipole-dipole interaction energy per atom and the 0.7 MHz/G ground state Zeeman splitting for ^{87}Rb[13] of ~300 pT (= 3 μG). At field strengths less than this the dipolar interaction adds further complexity to the dynamical evolution of the BEC and enables the study of magnetism at a quantum level.

3

Experiment

The apparatus constructed to realise our proposed experiments centrally consists of a long vacuum chamber with a small cross-sectional area. This is compartmentalised into a high-vacuum 2D MOT chamber where our atoms are pre-cooled and an ultra-high-vacuum science cell where our experiments will take place. These two regions are separated by a Differential Pumping Stage (DPS) to prevent the 2D MOT chamber, containing a high vapour pressure of rubidium, from contaminating the science cell. Vacuum is maintained within the chamber by use of a small ion pump midway along the DPS and a larger combined ion and getter pump attached to the science cell.

Atoms launched from the 2D MOT are recaptured at the extreme end of the

science cell in a 3D MOT where they are further cooled. Here they will undergo a stage of optical molasses cooling to assist loading into an optical dipole trap where they will be evaporatively cooled to BEC. This all-optical preparation of the BEC leaves the spin of the atoms as a degree of freedom as no cooling or trapping mechanism is spin-dependent, allowing for rich dynamical behaviour of the resulting BEC.

In considering the design of our vacuum chamber we will first address the aspects which will be common to many atomic physics experiments before covering the additional requirements specific to our experiment. The experiment-specific aspects of our chamber largely amount to the details of the glass cell where our experiments will be performed following perparation in the main vacuum chamber. Whilst we cannot ignore the fact that we will be using such a cell during the preceeding discussion we will only describe it in terms of its geometry until it is the central focus of discussion.

Vacuum 3.1

In order to produce a system with sufficiently weak interactions to reach BEC we are compelled to prepare a dilute gas of our chosen atomic species in a vacuum chamber. The density of the gas we aim to prepare is on the order of 10^{12}–10^{14} atoms/cm^3 which is several orders of magnitude less dense than a typical gas (around 10^{18} atoms/cm^3), leading to a very large mean interparticle spacing and thus little interaction between atoms. Alkali metals are typically used in atomic physics experiments due to their relatively simple energy level structure which contains the required "closed" transitions to utilise laser cooling. Rubidium (Rb) has been extensively used throughout the development of the atomic physics field and for this reason much of the requisite equipment is highly advanced and comparatively low cost, allowing for experiments based on Rb to be set up with relative ease. Such

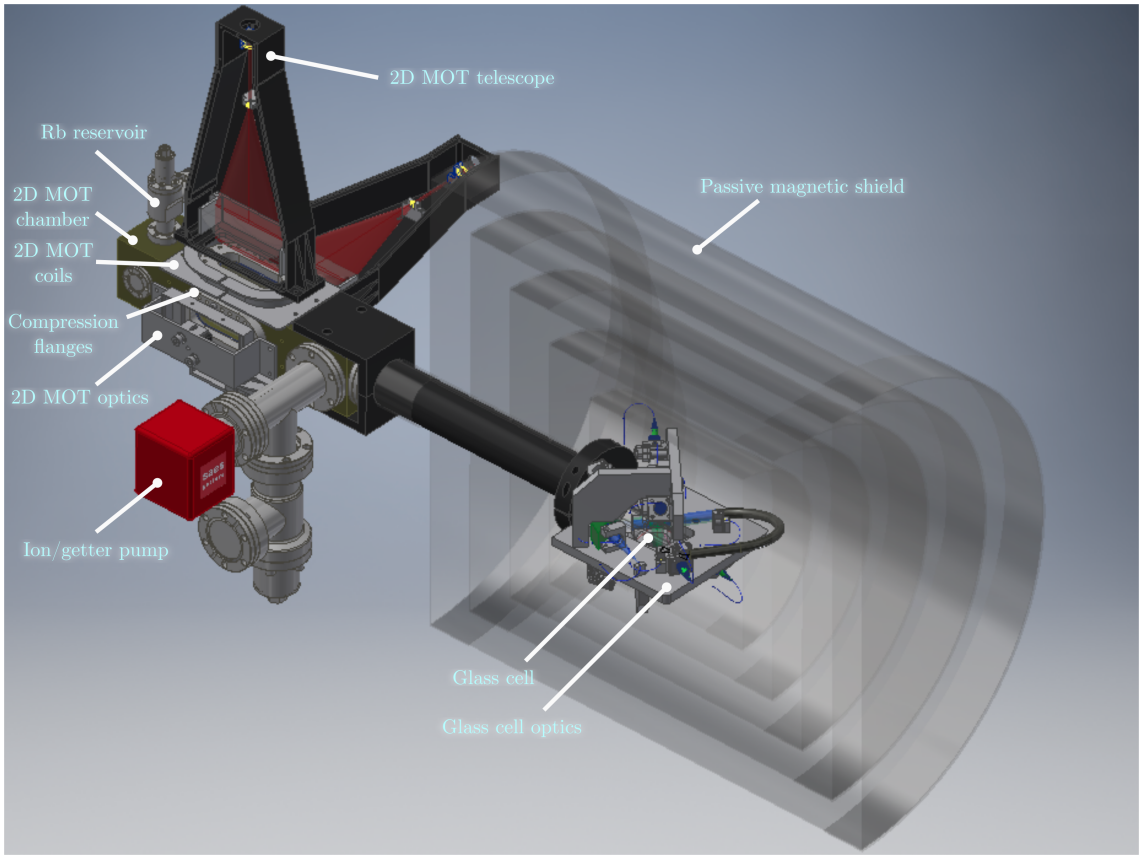

Figure 3.1 Our overall experimental setup with several key components labelled. We will discuss the Rb reservoir in section 3.1.2, the 2D MOT windows, optics, telescopes, and coils in sections 3.1.3–3.1.6, the glass cell in sections 3.1.10–3.1.11, and the passive magnetic shield in sections 3.3–3.3.4.

experiments are undertaken in vacuum in order to cause Rb (a solid at standard temperature and pressure) to exist in a vapour phase. The vapour pressure of Rb is around 5×10^{-7} mbar at room temperature[13], meaning that when the ambient pressure is reduced to this level the energy barrier to the vapour phase is reduced to zero. Lowering the ambient pressure further makes the vapour phase energetically favourable and thus statistically more likely to exist than the solid phase as once the atoms move apart they are unlikely to reach a situation where large numbers of atoms can collide in order to form a solid phase. The range of pressure from 10^{-7}–10^{-9} mbar is commonly referred to as High Vacuum (HV) and it is clear that we need at least this pressure everywhere in the chamber in order to thermodynamically support our vapour Rb phase. However at these pressures the collisional cross-section for 3-body collisions is still sufficiently large to make the phase space density required to cross the BEC transition very difficult to reach and would cause a BEC to decohere very rapidly due to collisions with background gas atoms. Short BEC lifetimes mean that any preparation and subsequent measurements of the BEC required to constitute an experiment must be extremely rapid, which can be very difficult to achieve.

In order to discourage 3-body collisions and to increase the lifetime of an achieved BEC we need to enter a regime of even lower pressure, 10^{-10}–10^{-12} mbar, commonly called Ultra High Vacuum (UHV). To reach these incredibly low pressures (15 orders of magnitude lower than atmospheric pressure) one needs to ensure that the UHV chamber is immaculately clean and made from non-porous materials with vapour pressures much lower than these levels in order to not limit the achieveable pressure due to material outgassing. Whilst HV is relatively straightforward to achieve, UHV presents a considerable increase in required care and effort and thus we would like to limit the size of this UHV region to be as small as possible. The benefits of this are twofold: firstly that there is a lower total surface area (proportional to levels of outgassing) of the UHV region, but secondly that the achieveable base pressure

from the action of an available pumping speed is proportional to the volume of the chamber being evacuated — thus lowering the volume gives us the best chance of being able to reach these low pressures.

From this brief review we can already begin to form an idea of what our vacuum chamber might look like: a HV section where the atomic vapour can be prepared and a separate UHV section where our BEC can be produced and an experiment performed. One concern of having two regions of different pressures is how to support ths pressure difference, or in other words how to keep our relatively "dirty" HV section from polluting our UHV section. Luckily this is a fairly standard practice and is achieved by having a long, thin tube connecting the two chambers. This creates a very low conductance between the two, making the probability of a particle travelling from one side to another very low. In turn, this means that the effective pumping speed of a vacuum pump attached to a chamber at one side of this tube is very different to that at the other end of the tube, leading to a differential pumping of the two chambers. Such a tube is often called a Differential Pumping Stage (DPS).

Pumps 3.1.1

Several stages of pumping are used to reach and maintain the vacuum within our chamber. In order of use, these involve a backing pump, turbo-molecular pump, ion pump and finally a getter pump. A backing pump reduces the pressure from ambient to roughly 1 mbar by mechanically extracting volumes of air within the chamber in a similar manner to that of a car engine, only attached to a closed volume for air intake. Under this modest vacuum we are now able to use more sophisticated and powerful pumping techiques such as turbo-molecular pumping. This uses a series of angled spinning blades, much like a jet engine, which spin at such a rate that gas molecules do not travel fast enough to be able to traverse a path back through the spinning blades and into the chamber leading to a net reduction of the pressure in

the chamber. This techique is very effective and can reduce the pressure to levels of 10^{-8} mbar and lower. The previous techniques described involve physically removing volumes of gas from the chamber and thus require that the chamber be open to atmosphere (through the pump exhaust ports). For practical and safety reasons this is not desirable as a long term solution. Fortunately, pumping techniques are available which can maintain such pressures and lower by removing contaminants from the gas phase. Such techniques can operate in a sealed chamber at an already low pressure, meaning that once a low enough pressure has been reached the chamber can be sealed and bulky mechanical pumps removed whilst these specialised pumps maintain vacuum within the chamber. Having reached a sufficiently low chamber pressure where we can utilise further "sealed-chamber" pumping techniques we turn to ion pumps to reduce the chamber pressure further. An ion pump comprises two electrical plates supporting a very high electrical potential between them which are situated in a strong magnetic field gradient. When an impurity atom passes between the plates it can become ionised due to the high electrical field present where the resulting charged constituents undergo precessional orbits due to the Lorentz force in the strong magnetic field which prevents them escaping the pump. Over time these ions are drawn toward the plates themselves due to attractive Coulomb forces acting on them. This technique can maintain pressures down to 10^{-10} mbar or lower. In our case, the stray fields from the magnets of the ion pump need to be considered in order to meet our stringent low field requirements. To this end we have opted to use combined ion/getter pumps manufactured by SAES getters which produce very low stray fields, on the order of only 6 G at the outer surface of the pump[1]. Additional pumping speed is provided from getter pumps, which are comprised simply of highly reactive materials with large surface area, which remove impurities from the surrounding vacuum when they collide with the getter surface due to the strong bond formed.

[1]SAES getters, personal correspondence.

3.1.2 Alkali source

A commonly used method to introduce alkali metal vapour into the vacuum chamber is to to place a small, partially-open container containing a mixture of a chromate salt of the chosen alkali and a thermally-activated reducing agent. Passing a current through a wire embedded in the container heats its contents and produces pure alkali vapour, and other gases, in the resulting redox reaction. An alkali salts have a lower the vapour pressure than the pure alkali metal, giving the experimenter control over the production of the alkali vapour at low pressures. In order to prevent the non-alkali product gases of this reaction from polluting the vacuum a reactive material, which can be the reducing material itself, is also included in the container. When gas molecules are incident on the surface of this highly reactive material they become strongly bound to its surface, effectively removing them from the residual gas within the chamber. Inclusion of such materials, commonly called "getters", is a common practice within vacuum chambers in order to achieve and passively maintain a given pressure. This overall device is called a dispenser and many types are available commercially[46]. Dispensers typically contain on the order of 10 mg of alkali metal which is expected to last a couple of years before needing replaced.

An alternative technique to produce alkali vapour in a vacuum chamber is to place a glass ampoule of pure alkali metal in a small volume of the chamber, a reservoir, which can be sealed-off with a valve. After the entire system has reached a satisfactory base pressure, the reservoir can be sealed and the glass ampoule therein broken causing the contained alkali metal to vapourise. The sealed reservoir now contains a very high partial pressure of alkali vapour which can be introduced in a controlled manner to the main chamber via the connecting valve in order to control the partial pressure of alkali vapour in the main chamber. Reservoirs can contain a much larger quantity of alkali metal than dispensers, typically on the order of 1 g, meaning that they are able to maintain a given partial pressure of alkali vapour

in the main chamber for a much longer time than a dispenser, avoiding the time-consuming and potentially risky[2] maintenance of venting the system to replace a spent alkali source. An additional benefit is that the pure alkali metal acts as a getter and can be used to adsorb chamber contaminants as it slowly coats the inner surfaces of the chamber.

In this experiment we have opted to use a Rb reservoir to produce our alkali vapour in order to prolong the maintenance-free lifetime of the chamber. This reservoir was realised as a 10 cm section of CF16 bellows connected to the main chamber via a right-angle valve and terminated on the other end by a blank flange. A glass ampoule containing 1 g of "natural" Rb (i.e. not isotopically purified) under an Argon atmosphere[3] was placed in the bellows alongside two M6 hex bolts in order to ensure the ampoule becomes readily trapped against a hard surface upon flexing the bellows. Once the chamber was been baked and reached a sufficient base pressure in the main chamber the reservoir valve was sealed and its bellows were flexed slightly beyond the point of trapping the ampoule several times in order to first score and damage the surface of the glass before finally trapping the ampoule and applying a sharp but controlled force on the bellows in order to break the glass. The break can be detected audibly and by the extended free-flexing range of the bellows. Initially the reservoir valve was opened halfway and the bellows heated to around 40 °C to aid evaporation of the Rb until a slight increase in base pressure was observed in the main chamber, whereupon the valve was closed again. In practice Rb vapour is slowly lost to the vacuum pumps and adsorbed onto the chamber walls, lowering the partial pressure and ultimately reducing the size of our MOT, requiring Rb vapour from the reservoir to be added periodically. In our case this is around once per week by slightly opening the reservoir valve for around 10 min.

With a controlled source of alkali vapour in our chamber we can now engineer

[2]Any vacuum maintenance involving venting the chamber or replacing parts inherently risks contaminating the internal surfaces of the chamber or damaging existing vacuum seals, thus it is ideal to design chambers to be as maintenance-free as possible.

[3]Purchased from Alfa Aesar.

ways to cool and trap the constituent atoms in order to perform experiments on them. In the HV section of our vacuum chamber we wish to produce a 2D MOT (i.e. atoms are cooled and trapped in 2 dimensions thus forming a high flux, low divergence atomic beam along the untrapped axis) in order to load the 3D MOT situated in the neighbouring UHV section of the vacuum chamber. This dual-MOT configuration allows us to reduce our experimental cycle time as the 2D MOT is able to load rapidly due to the high background pressure of Rb whereas we wish to have the science cell at a low background pressure in order to have a long lifetime of our BEC. By pre-cooling atoms in a 2D MOT and transferring these to a 3D MOT in the science cell allows us to have both long BEC lifetimes as well as short preparation times. To form the 2D MOT we require 2 counter-propagating beam pairs, to provide laser cooling, and a quadrupole magnetic field, to provide trapping. We will deal with these requisite elements separately and first focus on the requirements of the beams.

2D MOT: Windows 3.1.3

To load cooled atoms from our 2D MOT into our 3D MOT we need to align the untrapped axis of the 2D MOT with the axis of our chamber. As we require counter-propagating beams along two axes to form the 2D MOT we need to engineer optical access on four faces of the chamber. This can be realised using commercially available viewports which have circular optical elements vacuum-brazed directly to a standard CF mounting flange with a mating void to house the optic. Due to the high temperatures required in this process low-melting point glasses, such as the commonly used borosilicate glass, can become distorted around the weld or potentially across the entire optic. This is an issue as we cannot control the polarisation of the beams if our optical elements produce distortions on the wavefronts larger than $\lambda/4$. Whilst borosilicate glasses are used for almost all standard optical elements and are available with extremely good flatness specifications (as low as $\lambda/20$),

due to the nature of the brazing process such properties are not guaranteed to be maintained after being installed in a flange and any area over which they apply will vary for individual pieces[4]. Optics made from high-temperature melting point glasses, such as quartz or sapphire, can be substituted to avoid this problem but at significantly higher cost.

Aside from optical flatness there are other physical considerations which discourage their use. The highest initial velocity of an atom that will become trapped by the MOT is called the capture velocity of the MOT. To increase the capture velocity of the 2D MOT, in turn increasing the numbers of atoms trapped and thus the flux of the resultant atomic beam, one wishes to use large diameter beams. For a given borosilicate viewport the usable area available may be much smaller than the diameter of the optic, limiting the achievable capture velocity. It has also been shown for 2D MOTs specifically that elongated beams (i.e. having a large aspect ratio in their cross-section) lead to a higher flux of the resulting atomic beam[47]. Clearly standard viewports, being circular in shape, are not ideally suited to passing beams of a high aspect ratio with a minor axis on the order of several mm.

From the above considerations we are drawn towards a solution which can provide viewports with a high aspect ratio and a large optically flat area. A commonly used method to provide an UHV-grade seal between custom parts in a vacuum chamber is indium sealing. Caveats relating to the technique, detailed in the next section, provide some guidance on the desired geometry of the window and of the mating face of the chamber. For reasons which will become clear in the following discussion we are led to choose a window geometry which is rectangular with highly rounded corners in order to avoid issues during the sealing process. Ultimately, our 2D MOT windows take the form shown in figure 3.4 and are Anti-Reflection (AR) coated for 780 nm light at normal incidence. This AR-coating not only provides more light power to the chamber but avoids any interference effects which could occur within

[4]Kurt J. Lesker Ltd., private correspondence.

the MOT due to stray reflected beams.

Indium sealing 3.1.4

Welding is typically used to join two metals, however traditional welding techniques are not suitable for joining dissimilar metals or those of largely differing melting points. By bringing two metal parts in close contact under sufficient applied pressure and temperature the surface atoms of each part can be encouraged to diffuse into one another. The resulting joint from this diffusion-bonding process has comparable strength to that of a welded joint and is able to be used in situations where conventional welding would not be possible[48].

Indium is a very soft metal and can be diffusion-bonded to other materials at room temperature under only modest applied pressure. These properties along with indium's UHV-compatability allow indium to be used as a sealant in vacuum applications, for example between a flat window and a flat mating face of a vacuum chamber. In our case indium allows us to use bespoke cut glass to form viewports with a relatively large aspect ratio but this technique is generally useful whenever vacuum equipment with a non-standard mating adapter is required. Larger windows allow us to produce a larger 2D MOT capture volume for Rb atoms and results in a high quality optical element over a larger area than a standard flange-mounted equivalent.

The pressure required in order to diffusion-bond indium to another material, P, varies with temperature, T, which can be seen in figure 3.2. In practice one needs to convert this pressure into a required force exerted upon the indium by

$$P = aT^2 + bT + c \tag{3.1}$$

$$\therefore F = A\left(aT^2 + bT + c\right), \tag{3.2}$$

where A is the surface area of the indium seal.

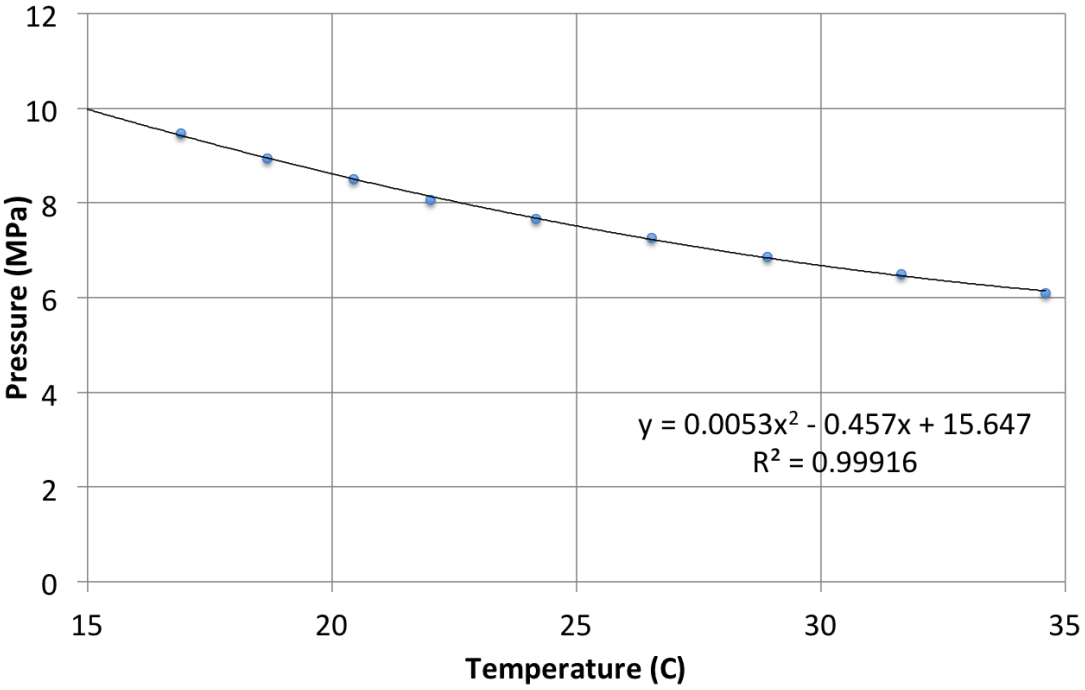

Figure 3.2 Minimum applied pressure to form a reliable indium diffusion bond against temperature when applied for 1 hr. Figure reproduced from a subset of the data presented in[48] (page 44) and fitted with a second order polynomial. The figure of merit used here, R^2, is a measure of the goodness of fit, where $R^2 = 1$ indicates a perfect fit of the model to the data. Here, this high value of R^2 gives us confidence that using this model of the required pressure for diffusion bonding with temperature well represents the actual data.

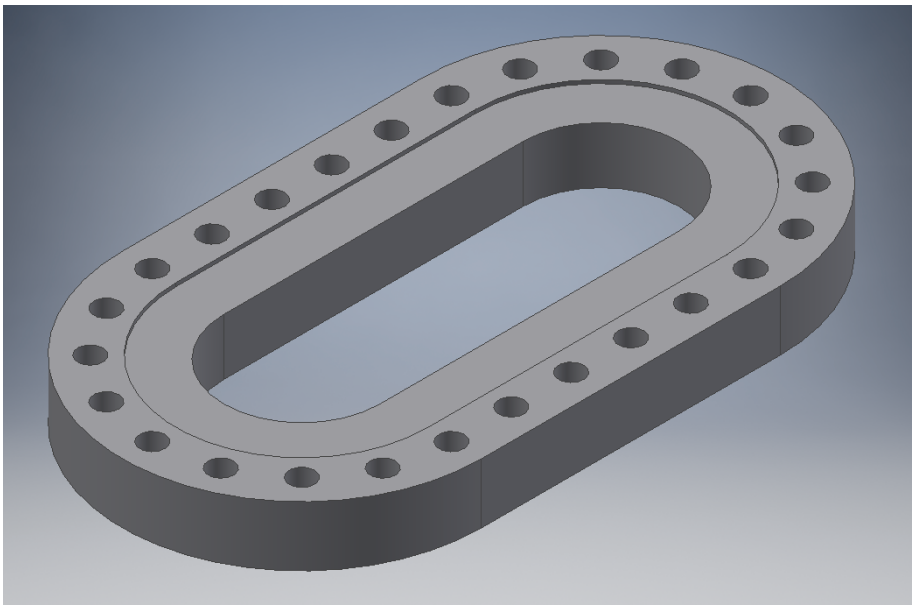

Figure 3.3 The compression flanges designed to apply force to the indium-sealed windows on the 2D MOT chamber. This view is the underside of the flange, which presses against the PTFE spacer. Note the recess to hold the spacer in place during assembly.

To apply force to the indium I designed compression flanges which were secured to the chamber with M4 bolts, shown in figure 3.3. When more bolts are used there is a reduced risk of applying uneven forces across the indium whilst sealing which can lead to the seal not adhering well, being squashed too thin to make a reliable seal, or sliding away from its original location underneath the compression flange. The heads of the mounting bolts are attached with a spring washer between two flat washers in order to enable a smooth loading of force onto the indium and to allow for small amounts of thermal expansion without damage to the glass or seals. The compression flanges are designed to match the shape of the windows and their profile has been taken into account when designing the forms for the required magnetic field coils which fit concentrically around the compression flanges so that they may remain in place.

Force is transferred from the compression flanges to the glass via a PTFE spacer (a soft plastic with very high melting temperature) in order to further spread the forces applied and prevent direct glass-to-metal contact with the compression flange.

As glass is such a hard material it is consequently very brittle and rapidly shatters under mild deformation. In order to ensure that the glass does not deform whilst tightening the compression flange the PTFE spacer is cut to match the shape of the indium seal such that forces are only applied over the area of the seal. In principle the compression flanges can be removed after the seals are made and the chamber is under vacuum in order to provide greater optical access. However, this is generally not recommended as indium has poor structural properties meaning that mechanical shocks to the chamber may cause the windows to shift and damage the underlying seals.

To ensure that the compression process is controlled we can use a torque wrench to apply a known torque to the retaining bolts. Using an empirical equation[49] we can calculate the axial force, F, from an applied torque, τ, on a bolt with a given thread diameter, d, and torque co-efficient, k

$$F = \frac{\tau}{kd}. \tag{3.3}$$

The torque co-efficient, k, depends on the materials of the bolt and thread as well as lubrication (which lowers the value of k) with typical values being $k \approx 0.2$. Substituting equation 3.3 into 3.2 and re-arranging gives

$$\therefore \tau = kdA\left(aT^2 + bT + c\right). \tag{3.4}$$

The shape of the indium seal used in this case is rectangular with two half-circles on either end, as shown in figure 3.4, which matches the shape of the glass to be fitted onto the chamber. Rounded edges are preferred when sealing with indium as an excess of indium is produced at sharp corners which can require excessive torque to flatten. This additional torque applies localised force risking breakage of the glass (catastrophic), and tearing of the indium giving rise to potential microleaks (very difficult to diagnose).

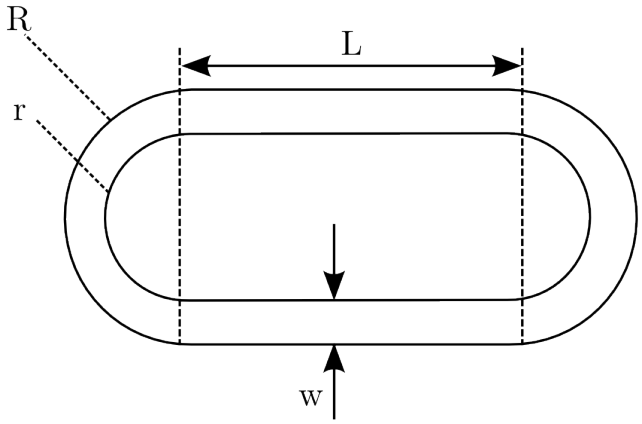

Figure 3.4 The shape of the PTFE spacer, 2D MOT windows, and indium seal. $R = 22.5$ mm, $L = 45$ mm, $r = 13$ mm, $w = 9.5$ mm, thus the dimensions of the 2D MOT windows are $2R + L \times 2R = 90 \times 45$ mm and are 6 mm thick.

When indium seals are made between two flat faces it is critical that these are both machined in a manner such that any residual scratches on the sealing surface are as small as possible and more critically aligned with the direction of the seal. This typically results in circular sealing faces, as these can be produced on a lathe with any residual scratches being parallel to the seal at all points. If a sealing surface has to be milled then the cutting bit will leave residual circular scratches along the length of the seal producing many scratches which are perpendicular to the seal. It can be difficult to deform the indium to fill these small scratches in order to seal them and thus form capillaries between the vacuum and outside environment and can lead to an unacceptably high leak rate of the seal. As the majority of our windows are non-circular a strategy to produce ideal sealing surfaces (i.e. where any surface scratches are all perpendicular to the seal) over faces which cannot be lathed had to be developed. To achieve this a small channel was machined along the length of the seal which is cut on a conventional mill. Whilst the tip of a milling bit produces circular scratches on the face underneath the bit the sides of the bit will produce only linear scratches along the walls of a channel milled in a material, producing ideal sealing surfaces of arbitrary shape. This small channel is

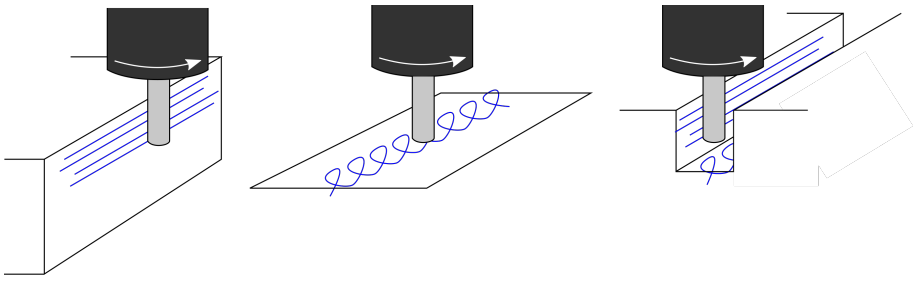

Figure 3.5 Illustrations of scratches left in a material due to milling (left) an edge, (centre) a surface, and (right) a channel.

also sufficiently deep to hold the indium in position whilst sealing, thus eliminating slippage during sealing as a failure mode.

To produce the seals a 1 mm diamter indium wire is placed in the machined channel on the chamber mating face, which is 1.1 mm wide and 0.5 mm deep. The indium wire has a greater volume than the channel in which it sits, by design, meaning that the indium should first fill and then spread out over the top of this channel when fully compressed. In our case, indium wire with a circular cross-section was used to create the seal which then deforms with applied force towards a T-shaped cross section (i.e. filling the channel and spilling over). For ease of calculation we will assume that the cross section is rectangular and constant with applied force. Here we will estimate that the width of the resulting indium seal is 1.4 mm (i.e. around 3× wider than the channel), where it can be shown that the surface area of the top of this indium seal is

$$A = \pi w^2 + 2w\left(\pi r + L\right). \tag{3.5}$$

It follows that the minimum required torque to form a seal, τ, is

$$\tau = kdw\left(\pi w + 2\left(\pi r + L\right)\right)\left(aT^2 + bT + c\right). \tag{3.6}$$

For the simpler case of a circular indium seal (i.e. $L \to 0$), such as the window at

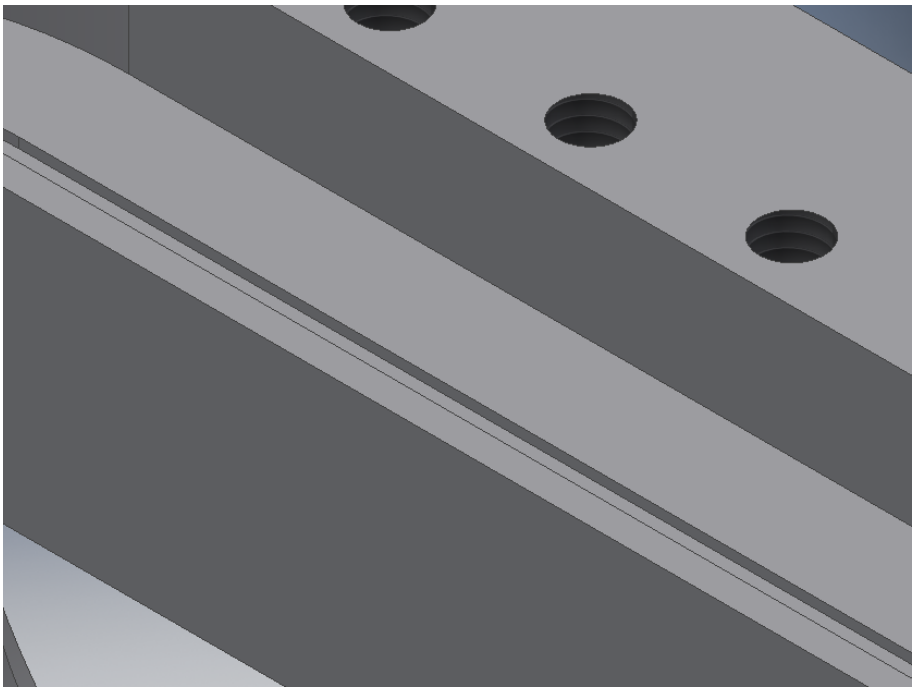

Figure 3.6 Close up view of the channels milled in the chamber to hold the indium wire and provide sealing faces. This can be seen as part of the larger vacuum chamber in figure 3.11.

the end of the chamber, this becomes

$$\tau = kd\pi w\left(w + 2r\right)\left(aT^2 + bT + c\right). \tag{3.7}$$

Using a value of $k = 0.22$ (based on a lubricated steel bolt in a titanium thread[49]) gives a $\tau \approx 3.4$ Nm for the elongated window and $\tau \approx 1.3$ Nm for the circular window at 20 °C. As this torque is the total produced by all of the compression flange bolts we should divide this value by the number of bolts in each flange to arrive at the torque required per bolt. For the elongated window this is $\approx 3.4/26 \approx 0.13$ Nm and $\approx 1.3/8 \approx 0.16$ Nm for the circular window. Both were tightened to 0.2 Nm per bolt in a cold room to ensure a good seal was made which was visually confirmed by the spreading of the compressed indium beyond the width of the sealing channel. It should be noted that 0.2 Nm of torque is *very* little and easy to exceed by hand, hence the care taken here to calculate the required sealing torque. One might intuitively

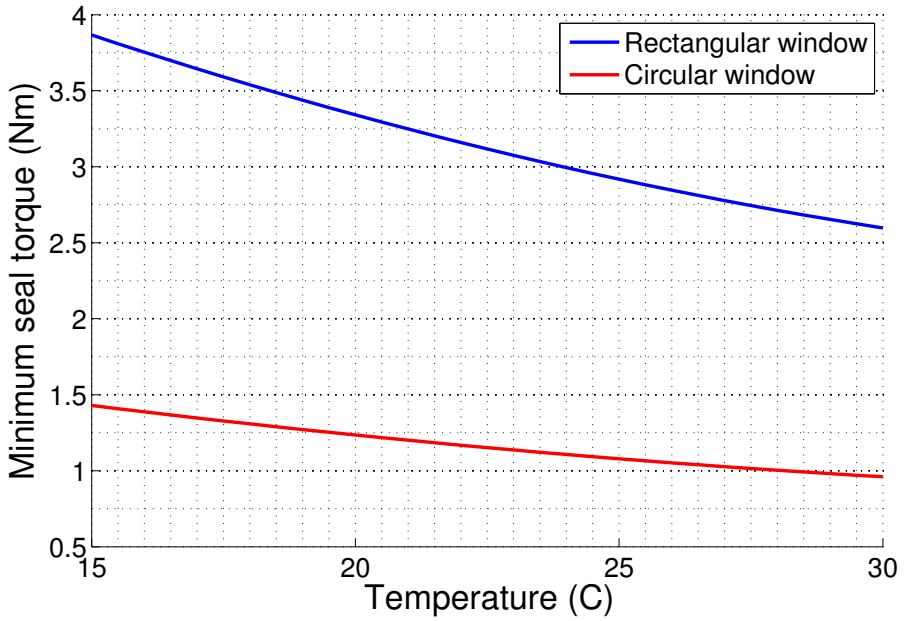

Figure 3.7 Minimum sealing torque vs temperature, calculated using equation 3.6 with $k = 0.22$.

think that over-tightening to something like 1 Nm should guarantee a seal under a wide variety of conditions, however over-tightening of the bolts places unneccessary strain on the glass and can squeeze the indium so thin that it no longer provides a sealing action between the faces.

To test that the seals were well made before continuing with assembly the chamber was connected to a helium leak detector with all other ports terminated in blank flanges. A helium leak detector comprises a small roughing pump, vacuum gauge and a very sensitive partial pressure detector for helium. These can be combined to determine the effective leak rate of helium based on typical atmospheric concentration and the pressure difference between the chamber and atmospheric pressure. Leaks can be further investigated globally by increasing the ambient helium content around the whole chamber or locally by directing a brief helium jet at an area of the chamber. Large leaks cause an almost immediate and dramatic spike in the effective leak rate, whilst delayed spikes can indicate that helium has diffused around the chamber to another leak nearby or that there is a capillary-like leak in that area.

In order to get a high degree of localisation of the leaks we connected a hypodermic needle to the end of a low pressure helium line which we could pulse by squeezing the soft material of the tubing to stop the flow and briefly releasing this to allow some helium to pass. After reapplying the same 0.2 Nm torque to some of the areas which showed spurious readings the entire chamber showed a global leak rate limited by the detector (1×10^{-12} mbar l/s, meaning that a chamber pressure of 1×10^{-12} mbar would give rise to an effective leak rate of 1 l/s for helium) indicating that the seals were adequately made.

2D MOT: Optics 3.1.5

An evanescent fibre-coupler network is used to split and mix our source beams in various combinations into the requisite beams, in-fibre, for our experiment. To realise the beams for the 2D MOT, which I have mentioned should have a cross-section with large aspect ratio, we must use a telescope to transform the output mode from an optical fibre. The telescope shapes the beam by first collimating the output of an optical fibre, then expanding this beam in one axis using a cylindrical lens and finally recollimating the expanded axis. In order to pass through the chamber without clipping of the beam we must constrain the final beam profile to fit within the viewing area through the indium seal compression flanges, shown in figure 3.3. The final recollimation stage within the telescope inadvertantly causes the non-expanded axis of the beam to be focussed, which is undesirable, however the focal length of the final lens is sufficiently long so as to not hinder the operation of the 2D MOT. This unwanted focussing could be avoided through the addition of a further cylindrical lens, however this is complicated due to the commercial availability of the requisite combination of compatible components. The telescope I designed, and the corresponding beam profile, are shown in figure 3.8.

In order to produce the correct handedness of circular polarisation for the retro-reflected beam in each axis we must use a $\lambda/4$ waveplate before a mirror. If we

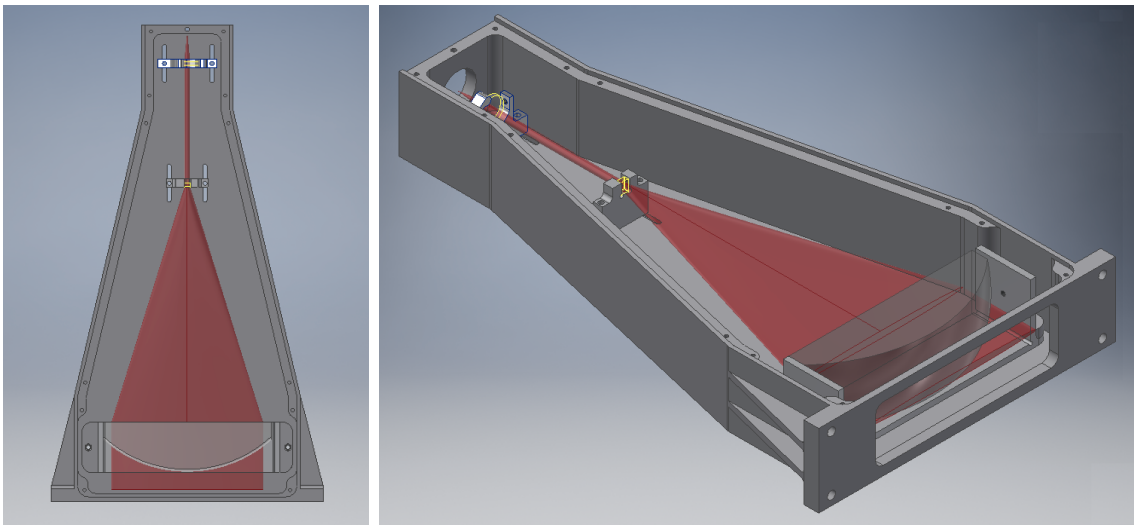

Figure 3.8 Telescope design used to shape the optical fibre output to that required for the 2D MOT (*left*) from above (*right*) isometric projection. Light enters the telescope via a fibre coupler, is collimated by the first lens (an $f = 25$ mm biconvex lens), expanded in one axis by the second (using an $f = 8$ mm concave cylindrical lens), then refocused by the final lens (a rectangular section of a larger spherical lens with $f = 150$ mm). The overall unit is around 40 mm tall, 150 mm wide, and 250 mm in length. Distances between the components have been chosen such that the light exiting the telescope is as collimated as possible. The completed unit is able to be mounted directly to the vacuum chamber.

did not use the $\lambda/4$ waveplate then we would have a reversal of the polarisation handedness from the point of view of the incident beam upon reflection. This means that the handedness of the polarisation from the point of view of the reflected beam (as per our chosen convention) is the same as that of the incident beam, which does not provide the correct configuration for atom trapping. Alternatively, including a $\lambda/4$ waveplate before the mirror first converts the incident circular polarisation to linear, which is preserved from the point of view of the incident beam upon reflection at the mirror, and is then converted back to a circular polarisation when passing through the waveplate again. This configuration preserves the handedness of the polarisation with respect to the incident beam upon reflection, as required for trapping.

The waveplate used[5] is optically flat ($\lambda/4$), multi-order, AR-coated for 780 nm light and is non-standard due to its large size, required in order to match the aspect

[5]Purchased from OptoCity.

ratio of the 2D MOT beams. This is mounted into a frame, seen in figure 3.9(c), using a strong epoxy glue which is in turn mounted to the retro-reflecting mirrors via nylon screws gripping the edges of the mirror. The waveplate mounting glue is applied by lying the frame and waveplate flat on a PTFE sheet before delicately applying epoxy with a toothpick to fill the gap between the optic and frame and allowed to set for 24 hr. Placing a PTFE sheet underneath the glueing process prevents the epoxy from strongly adhering to the underlying material or bench.

The mirrors used[6] are silver-backed with an optical flatness of $\lambda/10$. To achieve an aspect ratio of the mirrors to match that of the 2D MOT beams a large 10×10 cm mirror was precision cut into two 5×10 cm mirrors. These were carefully glued to low-profile kinematic mounts held in a custom frame, seen in figure 3.9(b), which is in turn mounted directly to the vacuum chabmer. This system provides two-axis adjustment of the retro-reflection mirror and waveplate with any vibrations in the optical system made common mode due to all being mounted directly to the vacuum chamber. The result is a highly stable optical system to support the production of our 2D MOT.

2D MOT: Magnetic fields 3.1.6

In addition to optical access we must also provide a means of producing magnetic fields in order to trap atoms. In our case we wish to produce a quadrupole field, which can be accomplished using a pair of coils with currents driven in opposite sense to one another. The separation between the coils and their geometry determine the steepness of the magnetic field gradient around the centre of their common axis. To maximise the field gradient for a given current the coils should be placed as close as possible to one another. We have chosen to mount these coils directly to the vacuum chamber, giving us the minimum possible separation for our geometry of 70 mm. The coils are wound onto "formers" which force the coil into a certain shape to allow

[6]Purchased from Edmund Optics and cut by Instrument Glass Ltd.

(a)

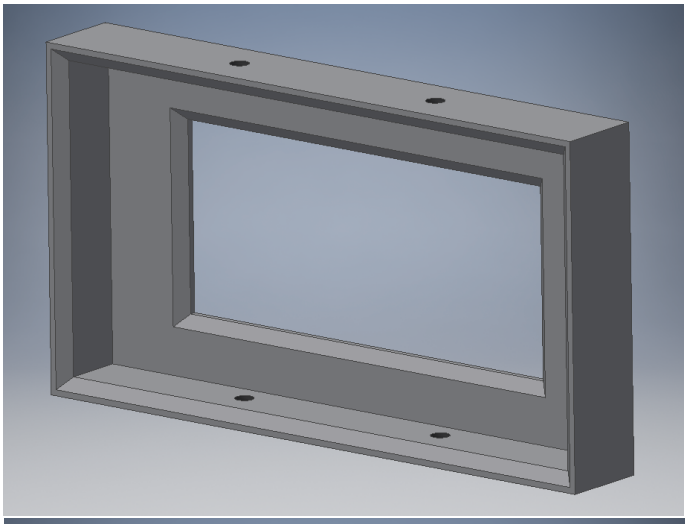

(b)

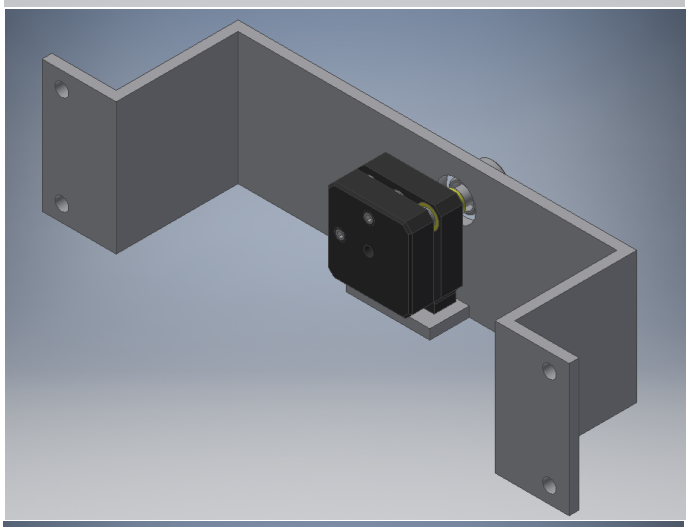

(c)

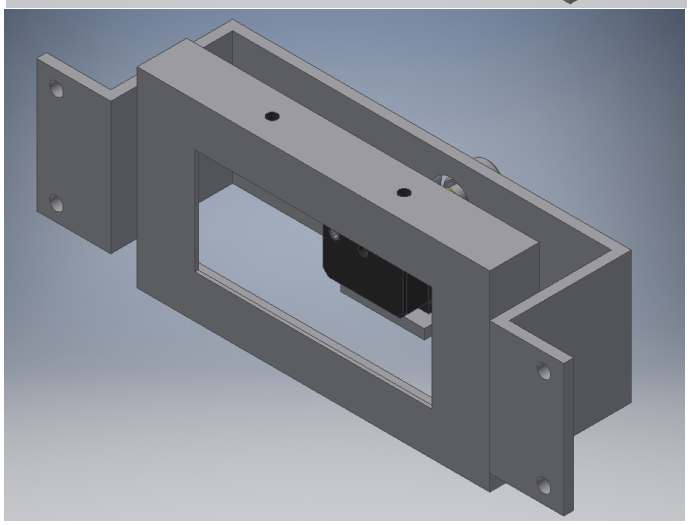

Figure 3.9 The assembly used to provide retroreflection of our 2D MOT beams. A waveplate is held in the frame seen in (a), which is mounted to a mirror by nylon screws which grip its edge. This mirror is glued to the kinematic mount seen in (b) and the whole assembly, seen in (c), is mounted to the vacuum chamber via the attached frame seen in (b).

for easier mounting of the entire assembly. Our formers must match the geometry of our compression flanges and thus our coils have an aspect ratio similar to that of our 2D MOT windows. 1 mm diameter wire was used to wind our coils as this will offer a much lower resistance than typical thinner gauge wires, leading to less heat dissipation during operation. Mounting directly to the chamber also allows some heat to be dissipated into the large thermal mass of the chamber body itself. It was observed that the temperature rise of the coils after an entire day's use was only around 20°C, thus requiring no additional cooling.

The field distribution produced by a coil can be computed by reducing the problem to the sum of a series of straight wires with a given current flowing through them. The resulting field distribution is then the sum of the field produced by each element. We can obtain the magnetic field ($\mathbf{B}$) produced at some distance ($\mathbf{R}$) from a wire (length $\mathbf{l}$ composed of incremental elements $\mathbf{dl}$) carrying a constant current (I) using the Biot-Savart Law

$$\mathbf{B} = \frac{I\mu_0}{4\pi} \int \frac{\mathbf{dl} \times \mathbf{R}}{|\mathbf{R}|^3}. \tag{3.8}$$

To more easily compute the resulting field for arbitrary geometries we can further generalise equation 3.8 in terms of points at the start and end of the wire segment and the point where we wish to evaluate the field. This will initially appear to complicate the expression however it ultimately allows us to separate the calculation of the magnitude of the field from the vector direction of the field. Consider a line in space, $\mathbf{l}$ (composed of incremental elements $\mathbf{dl}$), starting at $\mathbf{a}$ and ending at $\mathbf{b}$. The location we wish to calculate the resulting field profile at is $\mathbf{x}$. Note that $\mathbf{x}$ need not be situated at the origin of our co-ordinate system. The distance between $\mathbf{x}$ and $\mathbf{dl}$

is $\mathbf{R}$. Parameterising the line $\mathbf{l}$ with p ranging from 0 to 1 gives

$$\mathbf{l} = \mathbf{a} + (\mathbf{b} - \mathbf{a})p, \tag{3.9}$$

$$\mathbf{dl} = \frac{\mathbf{dl}}{dp}dp = (\mathbf{b} - \mathbf{a})dp, \tag{3.10}$$

$$\mathbf{R} = (\mathbf{x} - \mathbf{l}) = (\mathbf{x} - \mathbf{a}) - (\mathbf{b} - \mathbf{a})p. \tag{3.11}$$

Substituting these new expressions into equation (3.8) we see that

$$\mathbf{B} = \frac{I\mu_0}{4\pi}\int_0^1 \frac{(\mathbf{b} - \mathbf{a}) \times ((\mathbf{x} - \mathbf{a}) - (\mathbf{b} - \mathbf{a})p)}{|\mathbf{R}|^3}dp. \tag{3.12}$$

Since the cross product of anything with itself is zero and $(\mathbf{b} - \mathbf{a}) \times (\mathbf{x} - \mathbf{a})$ does not depend on p, we can write

$$\mathbf{B} = \frac{I\mu_0}{4\pi}\left[(\mathbf{b} - \mathbf{a}) \times (\mathbf{x} - \mathbf{a})\right]\int_0^1 \frac{dp}{|\mathbf{R}|^3}, \tag{3.13}$$

$$|\mathbf{R}| = |(\mathbf{x} - \mathbf{a}) - (\mathbf{b} - \mathbf{a})p| \tag{3.14}$$

$$= \sqrt{(\mathbf{x} - \mathbf{a})^2 - 2(\mathbf{x} - \mathbf{a}) \times (\mathbf{b} - \mathbf{a})p + (\mathbf{b} - \mathbf{a})^2p^2}. \tag{3.15}$$

Introducing the variables

$$\alpha = (\mathbf{x} - \mathbf{a})^2, \tag{3.16}$$

$$\beta = -2(\mathbf{x} - \mathbf{a}) \times (\mathbf{b} - \mathbf{a}), \tag{3.17}$$

$$\gamma = (\mathbf{b} - \mathbf{a})^2, \tag{3.18}$$

we can evaluate the integral in equation 3.13 as

$$\int_0^1 \frac{dp}{|\mathbf{R}|^3} = \int_0^1 \frac{dp}{(\alpha+\beta p+\gamma p^2)^{\frac{3}{2}}} \tag{3.19}$$

$$= \left[\frac{-2(\beta+2\gamma p)}{(\beta^2-4\alpha\gamma)\sqrt{\alpha+p(\beta+\gamma p)}}\right]_0^1 \tag{3.20}$$

$$= \left[\frac{-2(\beta+2\gamma)}{(\beta^2-4\alpha\gamma)\sqrt{\alpha+(\beta+\gamma)}}\right] - \left[\frac{-2\beta}{(\beta^2-4\alpha\gamma)\sqrt{\alpha}}\right] \tag{3.21}$$

$$= \frac{2\left(\beta\sqrt{\alpha+\beta+\gamma}-(\beta+2\gamma)\sqrt{\alpha}\right)}{(\beta^2-4\alpha\gamma)\sqrt{\alpha^2+\alpha(\beta+\gamma)}}. \tag{3.22}$$

Substituting this back into equation 3.13 we obtain

$$\mathbf{B} = \frac{I\mu_0}{4\pi}\left[(\mathbf{b}-\mathbf{a})\times(\mathbf{x}-\mathbf{a})\right]\left[\frac{2\left(\beta\sqrt{\alpha+\beta+\gamma}-(\beta+2\gamma)\sqrt{\alpha}\right)}{(\beta^2-4\alpha\gamma)\sqrt{\alpha^2+\alpha(\beta+\gamma)}}\right], \tag{3.23}$$

$$\therefore \mathbf{B} = \underbrace{\left[\frac{I\mu_0\left(\beta\sqrt{\alpha+\beta+\gamma}-(\beta+2\gamma)\sqrt{\alpha}\right)}{2\pi(\beta^2-4\alpha\gamma)\sqrt{\alpha^2+\alpha(\beta+\gamma)}}\right]}_{\text{Scalar part}} \underbrace{\left[(\mathbf{b}-\mathbf{a})\times(\mathbf{x}-\mathbf{a})\right]}_{\text{Vector part}}, \tag{3.24}$$

where the field magnitude and direction calculations have been largely separated.

For simplicity we modelled our coils as rectangular rather than their true geometry. This will lead to some error between the calculated and fields produced from these coils but these differences should be within 20 % and allow us to get order-of-magnitude estimates of operating currents required. Using approximate coil dimensions of 100×60 mm separated by 77 mm leads to a field gradient at the midpoint between the coils along their axis to be 2.8 G/cm/A. Typically a MOT requires a gradient of around 15 G/cm[7], meaning we will need to use currents around 5 A. This amperage is within the limits suggested based on the gauge of the wires used to wind the coils and can be readily supplied by standard bench power supplies. Due to the large wire diameter used in the winding of the coil and the heatsinking provided by the metal former and chamber body the temperature of the coil itself does not rise significantly with use over the course of a day. The temperature of the

[7]From the experience of other experiments within our group.

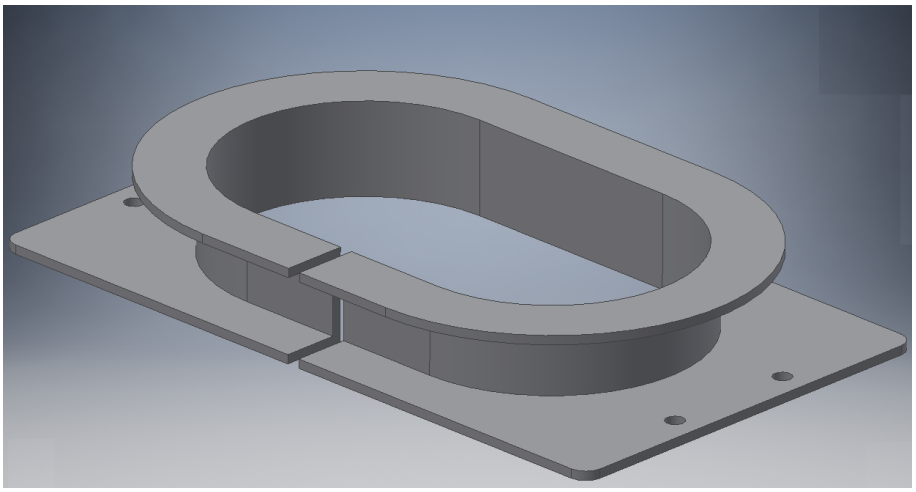

Figure 3.10 The formers used to wind our magnetic coils for our 2D MOT. These are sized to fit concentrically around the compression flanges for the 2D MOT windows (seen in figure 3.3).

coils were found to rise by around 20°C and the chamber did not noticeably increase in temperature.

2D–3D MOT population transfer 3.1.7

After producing our 2D MOT we must then transfer these pre-cooled atoms to our 3D MOT located in the adjacent UHV chamber. This transfer can be realised by physically moving the region in which there are trapping forces, for example using moving coils or by shifting the focal point of a high power laser using a mechanised telescope. In our case we opted to achieve this transfer ballistically, avoiding the complexities of large moving parts. 2D MOT atoms are accelerated in the direction of the 3D MOT (i.e. along the axis of the vacuum chamber) using an unbalanced cooling beam along the axis of the chamber. Being unbalanced, this beam provides a momentum kick to the atoms in a defined direction which is persistent and cumulative over many cycles. Fortunately, due to the Doppler effect, accelerated atoms eventually reach velocities which cause them to see this "pushing" beam with a sufficiently large detuning (and thus a proportionally low scattering rate) to be considered as not interacting with the pushing beam.

(a)

(b)

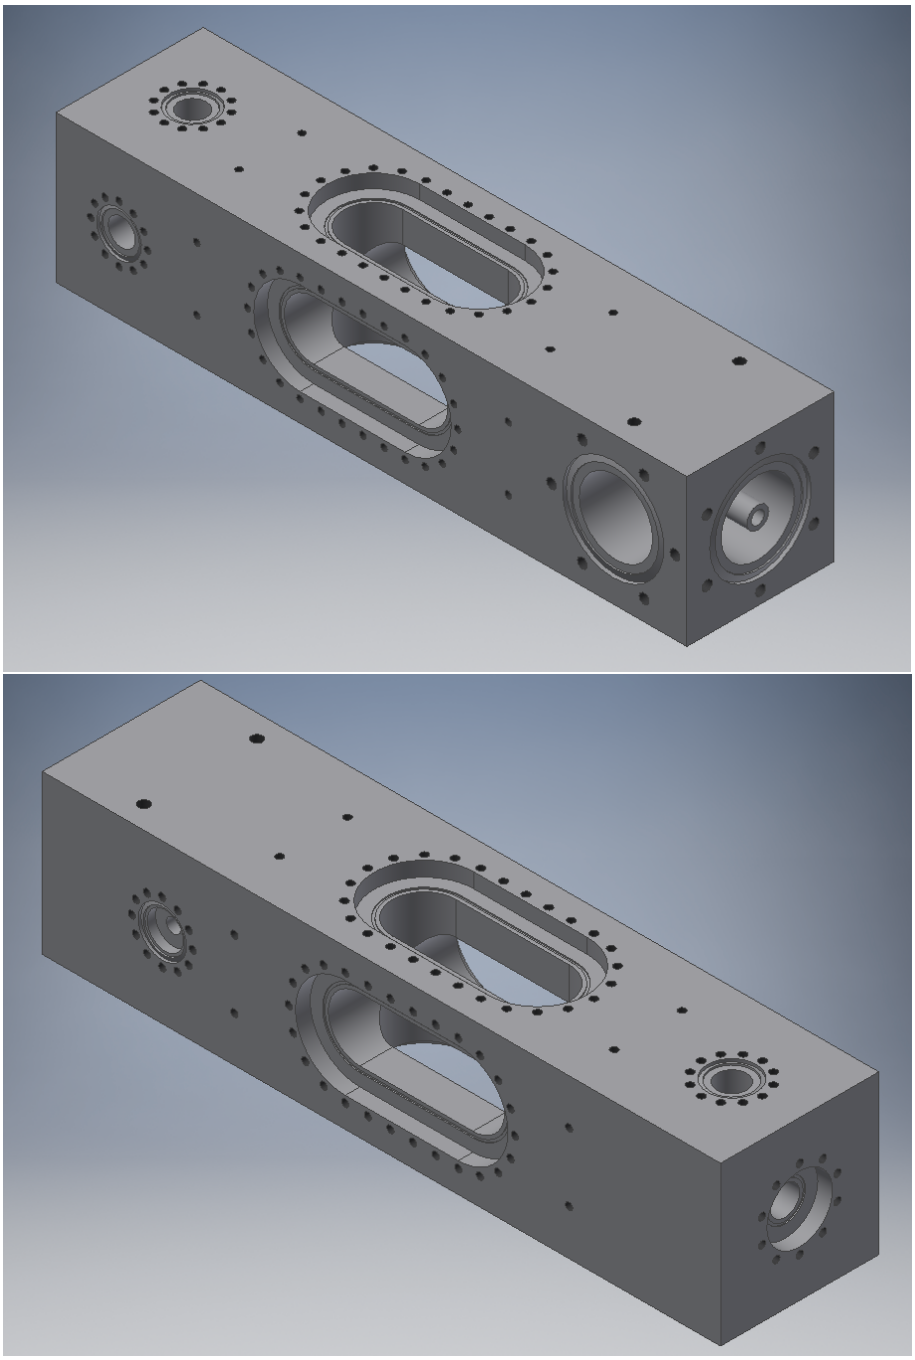

Figure 3.11 The main body of the vacuum chamber (a) looking from the UHV side of the chamber and (b) from the HV side.

The approach of a ballistic launch offers several advantages besides reduced mechanical complexity. Compared to an "unpushed" 2D MOT, this arrangement provides a lower divergence of the resulting atomic beam due to the increased axial velocity. This is crucial to enable large numbers of atoms to pass though the DPS without collisions with the internal walls, whereupon atoms would be adsorbed and thus lost. This ultimately leads to a higher atomic flux at the 3D MOT, allowing us to capture a greater number of atoms and giving us a better starting point to approach BEC.

Due to the high degree of localisation of the atoms due to their low temperature one can well-estimate their trajectory using simple Newtonian mechanics. For a given launch velocity and angle one can predict how far the atoms will have fallen under gravity on their path towards the 3D MOT. Clearly when this fall is comparable to the size of the 3D MOT we will suffer poor transfer efficiency between the 2D and 3D MOTs. In the simplest case where we chose a launch angle of 0° (i.e. atoms launched parallel to the chamber axis) we are compelled to launch the atoms with as high a velocity as possible in order that the atoms fall under gravity by the least amount over their trajectory. In practice, we are limited in how much we can increase the 2D MOT atoms' axial velocity due to the fact that we have an upper limit on the velocity of atoms which we can capture with the 3D MOT. In order to compensate for this we can increase the launch angle, allowing us to use a lower launch velocity to achieve a given fall distance over the trajectory of the atoms. Further, we must ensure that atoms' trajectory passes comfortably through the DPS. Assuming that our 3D MOT is spherical and roughly 2 mm in diameter this means we wish to constrain the total vertical displacement of the atoms after the launch to within ±1 mm of the chamber axis. Vertical displacement at the 3D MOT due to this ballistic launch versus launch velocity and angle, limited to ±1 mm, is plotted in figure 3.12. The upper limit of achieveable capture velocity for a MOT is around 50 ms^{-1}, hence we have limited the plot to this range. Selecting an angle between

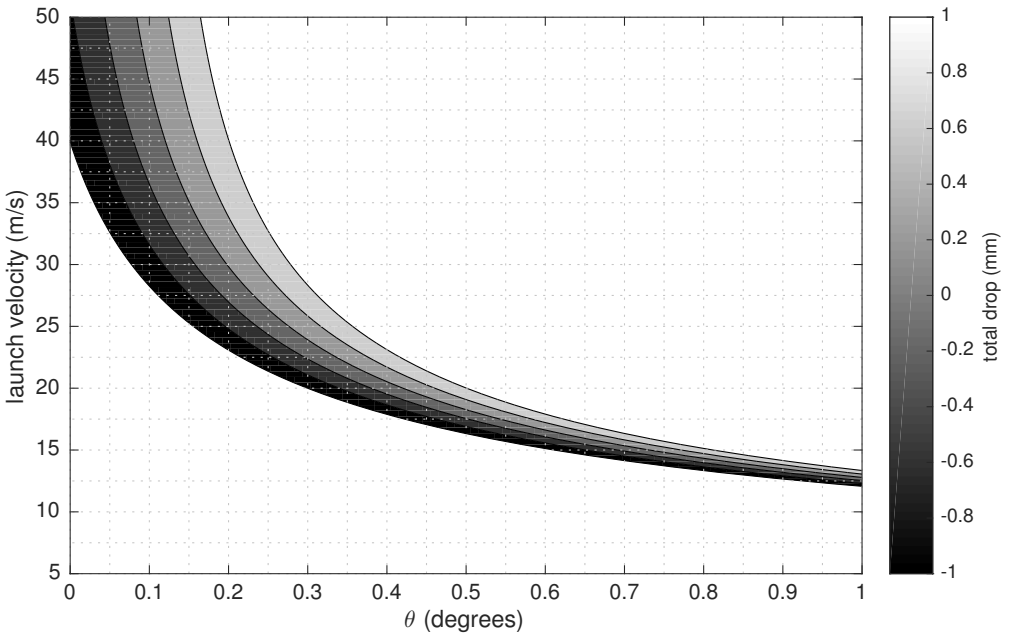

Figure 3.12 Resulting vertical centre-of-mass displacement following ballistic launch of atoms from our 2D MOT to 3D MOT versus launch angle and velocity.

0.1–0.3° allows us to keep the launch velocity within the range of capture of the 3D MOT whilst retaining a good tunability over the fall of the atoms. To further increase the atomic flux we could increase the power of the pushing beam such that we enter a regime of "guiding" (via the resulting dipole force) as opposed to pushing the atoms. Whilst launch velocity increases with power in the low power limit, as the power is increased we can produce dipole trapping of the atoms due to their being drawn to regions of high intensity. This can be used to essentially eliminate divergence of the atomic beam by providing radial trapping along the beam.

HV-UHV separation 3.1.8

To maintain a pressure difference between two sections of a vacuum chamber the connection between them must have a low conductance. Such a connection is often called a Differential Pumping Stage (DPS). In our case we opted to realise a low conductance line using a graphite tube, which additionally acts as a getter. As our

chamber is pumped by an ion pump located midway along the DPS a hole was carefully drilled perpendicular to the tube's axis to provide a path to the ion pump.

The unbalanced force acting on the DPS due to the pressure difference between is on the order of 7 pN under normal operation (i.e. $\Delta P \approx 10^{-8} - 10^{-11} \approx 10^{-8}$ mbar $= 10^{-6}$ Pa) however if the chamber were vented rapidly from the UHV side (i.e. $\Delta P \approx 10^{3} - 10^{-8} \approx 10^{3}$ mbar $= 10^{5}$ Pa) this could rise to 0.7 N which is large enough to potentially dislodge the DPS. To avoid the DPS becoming dislodged in any event a small wedge was used to retain the graphite tube within the larger tube machined into the chamber body itself. This wedge was made from a small section of UHV-grade oxygen-free copper which was sharpened to a point using a pair of side-cutters. With the graphite tube held in place by hand this copper "nail" was gently tapped into the small gap between the graphite tube and the chamber. For ease of installation, this wedge was placed on the UHV side of the DPS. The rigidity of the resulting assembly was gauged by touch to be satisfactory before attaching the glass cell.

3.1.9 UHV chamber

In order to achieve the highest effective pumping speed at the experiment region of the glass cell we must maximise all conductances between pumps connected to this region. Assuming that we cannot bring the pumps any closer, this means that we must maximise the cross-sectional area at each point along the path to the glass cell. We can achieve the greatest benefit from applying this at the junction connecting the main pump for the chamber and the opening of the glass cell as all other series conductances are as large as can be within reason. This region has been opened up as much as reasonably possible in order to provide the highest conductance through this section.

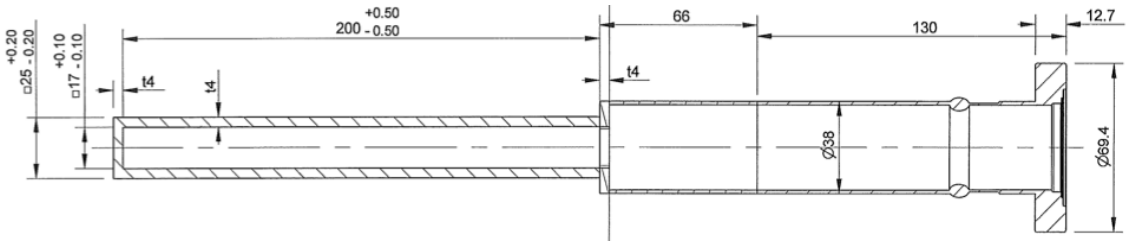

Figure 3.13 The working drawing used to construct our glass cell (all dimensions in mm). The section on the right is a cylindrical piece which forms the glass-to-metal transition to a CF35 vacuum flange, and the section to the left is cuboidal and formed of 4 rectangular glass plates which have anti-reflection coatings.

3.1.10 Glass cell: design

The glass cell used in our experiment needs to meet a multitude of requirements. The cell should be sufficiently long in order to allow for axial nesting of the passive magnetic shields used to reach the low field levels required for our experiment. In order to connect to the rest of the vacuum system the cell must have a glass-to-metal seal in order to be terminated in a standard vacuum flange. The primary requirement for our experiment is that the cell is conducive to an ultralow magnetic field environment, which means that the glass-to-metal seal should be made using a non-magnetic metal (such as 304 stainless steel) and the glass must have a very low level of magnetic impurities. A design based on these requirements can be seen in figure 3.13.

One can estimate the magnetic field produced from all of these magnetic impurity ions by first calculating the magnetic field at the experiment location due to a single ion within the glass. Considering Fe as the main magnetic impurity and taking the value of the magnetic moment of such an impurity to be $\mathbf{m} = 5\mu_B$ (roughly that of Fe^{2+} and Fe^{3+} [50]) we can estimate the magnetic field at a distance from this dipole

using

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi}\left(\frac{3\mathbf{r}(\mathbf{m}\cdot\mathbf{r})}{|\mathbf{r}|^5} - \frac{\mathbf{m}}{|\mathbf{r}|^3}\right) \tag{3.25}$$

$$= \frac{\mu_0}{4\pi|\mathbf{r}|^3}\left(3\hat{\mathbf{r}}(\mathbf{m}\cdot\hat{\mathbf{r}}) - \mathbf{m}\right). \tag{3.26}$$

The maximum field from the dipole is seen when $\mathbf{m} \parallel \mathbf{r}$, thus we will assume this in our calculations. With this assumption we can simplify the above equation to

$$\mathbf{B}(\mathbf{r}) = \frac{|\mathbf{m}|\mu_0}{2\pi|\mathbf{r}|^3}. \tag{3.27}$$

The internal cross-section of the glass cell is square with a side length of 17 mm bounded by 4 mm thick glass walls. We will choose to situate our impurity dipole at a point midway into the bulk of the glass however in practice these will be distributed throughout the glass. This means that a point in the centre of the cell's cross section, where the experiment will take place, is 10.5 mm from a point midway into any of the surrounding glass walls. This yields a flux density at the centre of the cell of roughly 6.4 aT per impurity.

Our glass cell is made from a synthetic quartz in order to have a low concentration of magnetic impurities. As will become apparent, this is critical to achieving a sufficiently low magnetic field for our experiments in the centre of the cell. The crystal structure of quartz has a unit cell containing 3 atoms with side length of roughly 5 Å, giving a density of atoms of roughly 2.4×10^{28} m^{-3}. Fe impurities are present in fused quartz at a level of roughly 5 ppm[51] however our synthetic quartz glass[8] has Fe impurity levels of 1 ppb[52], giving a density of roughly 2.4×10^{19} m^{-3}. Assuming that contributions from glass sections further away than the immediate walls (i.e. roughly within 2 cm^3) are negligible, we can treat the cell as 5 square pieces of glass assembled in a cube with an open end. Taking such a wall of the cell to be $17\times17\times4$ mm^3 then in 5 of such walls (comprising the cell with access from 2D MOT)

[8]Suprasil-P20, manufactured jointly by ShinEtsu Quartz and Heraeus.

we have roughly 1.38×10^{14} impurities which contribute to the overall field. Assuming that these are diluted enough within the bulk to not influence each other, we will not have all of these dipoles aligned with each other (i.e. ferromagnetic, $\mathbf{m} \propto N$) but a random relative alignment of the dipoles (i.e. $\mathbf{m} \propto \sqrt{N}$), yielding a net flux density at the experiment of $\sqrt{1.38 \times 10^{14}}\times6.4$ aT $= 75$ pT. As required, this results in a flux density lower than that required to discern dipole-dipole interactions (calculated in Section 2.5.3 to be roughly 300 pT) and highlights the importance of using synthetic quartz versus fused quartz. As fused quartz has a roughly 1000-fold higher impurity concentration this would lead to a roughly 30 times higher flux density at the centre of the cell, which would be unacceptable for our experiments.

Further to material considerations, we are also concerned with the optical properties of the glass. We require that the cell have very low reflectivity for 780 nm and 1550 nm light at various angles of incidence for both s- and p-polarisations. For the lowest reflectivity one must AR-coat both the inside and outside faces of the cell such that reflections at all interfaces are minimised. In order to reduce the complexity of the AR-coating on the cell, which in turn results in a higher performance coating in general, we wish to minimise the number of faces which require multi-wavelength coatings and reduce the number of incidence angles for a given wavelength on a given face. By a clever arrangement of the beams for our 3D optical lattice we are able to send all of its constituent beams through only one axis of the cell (i.e. two faces) with a single angle of incidence which greatly simplifies the AR-coating specification.

The resulting optical arrangement is that one of the 3D MOT beam pairs (780 nm) passes with normal incidence through the top-bottom faces of the cell and the other two beam pairs pass through the side-side faces of the cell with a 45° angle of incidence such that all beam pairs are mutually orthogonal. The optical lattice is formed by passing two orthogonal retroreflected beam pairs through the side-side faces of the cell with a 45° angle of incidence, forming a 2D lattice, which is overlaid

with two beams crossed at 90°, with a 45° angle of incidence to the cell face, in an orthogonal plane to that of the 2D lattice beams. A further pair of beams crossed at an angle (i.e. not directly retroreflected) produce a 1D lattice with a lattice vector derived from the difference in k-vector of the constituent beams. Another 1550 nm beam is required to realise the dipole trap, however its angle relative to all other beams is arbitrary and thus was chosen based on the lowest reflection co-efficient achieved by the AR-coating, calculated to be 60° passing through the side-side faces of the cell.

Glass cell: construction 3.1.11

To achieve the lowest reflection co-efficient for a beam passing through the cell it is desirable to AR-coat the glass on both sides, however this cannot be achieved for all construction methods. One method of glass bonding is frit-glass bonding, which acts much like a solder for glass and allows the creation of a solid glass join between separate pieces of glass. When joins between glass pieces are made with this method the resulting cell is effectively constructed from a single piece of glass is thus both structurally very strong as well as having extremely low outgassing rates. Due to the high temperatures used in this bonding process (> 1000°C) AR-coatings must be applied afterwards as they would not withstand these high temperatures (degradation or failure occurs in the region of 250–300°C). This means that a having a frit-glass bonded cell precludes the use of AR-coating both sides of the glass as the inside faces of a cell cannot be coated once assembled[9]. An alternative contruction method is to use optical contact bonding to join the plates. In this method the joining faces are machined to optical flatness (circa. $\lambda/10$ or better) and then carefully brought into contact and pressure applied until the pieces are joined via the Van der Waals' forces between them. This method also provides a

[9] AR-coatings are commonly applied by vapour deposition in vacuum such that the top-side of all elements are coated. This means that AR-coatings cannot be applied to the inside faces of assembled cells.

very low outgassing co-efficient but with reduced mechanical strength. With careful design of the glass plates this method can allow for AR-coatings on both sides of the glass before assembly. The downside to this method is its highly technical nature, meaning that few companies can offer this service which is often at high cost.

One may also use an UHV-rated epoxy (i.e. with a sufficiently low outgassing rate once cured) to construct a UHV cell. A thin layer of this epoxy is applied at the joins of the cell and is baked according to the specifications of the particular epoxy used in order to reach a fully cured state. With this method it is much easier to AR-coat both sides of the glass before assembly and the relatively simple nature of the process means that it is readily offered by manufacturers of glass cells. For these reasons we opted to use an epoxy-bonded glass cell (with internal and external AR-coating) in the first instance, however for reasons that will become clear this was later replaced with a frit-glass bonded equivalent (with external AR-coating only).

Using epoxy in a vacuum application is always a risk and should be considered as such. In particular, one should understand the nature of the epoxy used and its properties under baking, both of which are influenced by the initial curing method. A particular point of note is that for any plastic-like material with an undefined melting point (such as epoxy) there is a so-called glass transition temperature, T_g, where the material begins to become soft. Clearly, baking an epoxy-sealed vacuum system around this temperature should be avoided as it could introduce microleaks or affect the hermeticity of the seal. In general, curing an epoxy at high temperatures raises the resultant T_g and offers a higher range of baking temperatures for the vacuum system.

Our original glass cell[10] was epoxy-sealed using `Epotek-353ND-T`[11] which was cured by holding the cell at 120 °C for 5 min. A distributor of Epotek epoxies[12] revealed that they observed increased leak rates in tests when baking epoxy-bonded

[10]Manufactured by JapanCell.
[11]Hiro Yazawa, JapanCell, personal correspondance.
[12]Rob Church, Kummer, personal correspondance.

cells around T_g, shown in figure 3.15. Curing information for `Epotek-353ND-T`, shown in figure 3.14, provides an insight into where the glass transition might lie for a given cure method. For the method used by JapanCell, this should yield $T_g \approx 110$ °C. Figure 3.15 clearly shows a spike in the leak rate around 125 °C, which suggests that one of the more optimal curing methods was used for these tests, which pushes T_g as high as possible (see the lower section of figure 3.14). The width of this leak rate peak is around 50 °C and centred on T_g. These should set some minimum limits to be placed around the expected T_g for a given cure method wherein baking should not be held in this region for an extended period to avoid introduction of microleaks or cause other issues with the hermeticity of the seal. More troubling is that figure 3.14 provides data for a minimum cure time of 30 min which leaves 4% of the epoxy uncured and thus available to outgas and cause other vacuum issues. One would expect this uncured percentage to increase for shorter cure times such as those chosen by JapanCell.

Colleagues[13] have informed us that baking at well below T_g or (rather unintuitively) well above T_g both produce satisfctory results, so long as the region around T_g is crossed as rapidly as possible (within material thermal stress limits). Unfortunately this information was received too late whereupon we observed unacceptably high pressures in our UHV chamber. At this stage it appeared clear that we must have introduced microleaks in the epoxy seals during baking as this was conducted around 130 °C, with no special haste to heat beyond the region around T_g. Our use of indium seals also limits the upper baking temperature to 150 °C due to the melting point of indium, leading to a very narrow acceptable temperature range within which to bake the chamber. It was ultimately decided that the trade-off between a lower risk of a high leak rate developing during baking against a higher reflection co-efficient of the cell was a worthwhile one and an alternative frit-glass bonded cell was obtained[14]. Upon changing the cell the pressure in the UHV chamber was found

[13]Dr. Paul Griffin, University of Strathclyde, private correspondence.
[14]Manufactured by Precision Glassblowers of Colorado.

EPO-TEK 353ND Tg versus Cure Cycle

Cure condition	Residual Exotherm (J / g)	Peak exothermic Temperature (C)	% NOT cured (% unreacted)	Calculated Tg (C)
80 C / 30 minutes	328.5	134.6	65	87
80 C / 1 hr	253	137.4	50.7	87
80 C / 2 hr	212.3	138.8	42.5	87
100 C / 30 minutes	83.1	155.2	16.6	103
100 C / 1 hr	86.14	155.5	17.3	103
100 C / 2 hr	51.94	155.1	10.4	103
120 C / 30 minutes	20.16	158.2	4	110
120 C / 1 hr	8.67	159	1.7	110
120 C / 2 hr	0.8	155.6	0.2	110
150 C / 30 minutes	N/A	N/A	N/A	120
150 C / 1 hr	N/A	N/A	N/A	119.4
150 C / 2 hr	N/A	N/A	N/A	116.4
80 C / 1 hr + 120 C / 1 hr	19.9	158.4	4	110
80 C / 1 hr + 150 C / 1 hr	N/A	N/A	N/A	129.4
80 C / 1 hr + 120 C / 1 hr + 150 C / 1 hr	N/A	N/A	N/A	128.8
150 C / 15 second hot plate cure	N/A	N/A	N/A	137.3
totally uncured exotherm	498.9	118		

- notice that additional time added to constant cure temp does not influence Tg so much
- notice that percentage of cure can be influenced heavily by additional time at constant cure temp
- notice that any cure temp less than 120 C, will result in incomplete cure (and thus can expect to have further post-cure stresses in the field or down-stream process)
- notice that 150 C / 15 sec hot plate cure gives highest Tg. This is INDUCTION cure vs. BOX OVEN
- percentage NOT CURED was determined by ratio [residual exotherm / 498.9 J-g x 100 %]

Figure 3.14 Internal data (provided by Rob Church of Kummer) on the expected glass transition temperature, T_g, for a given cure method.

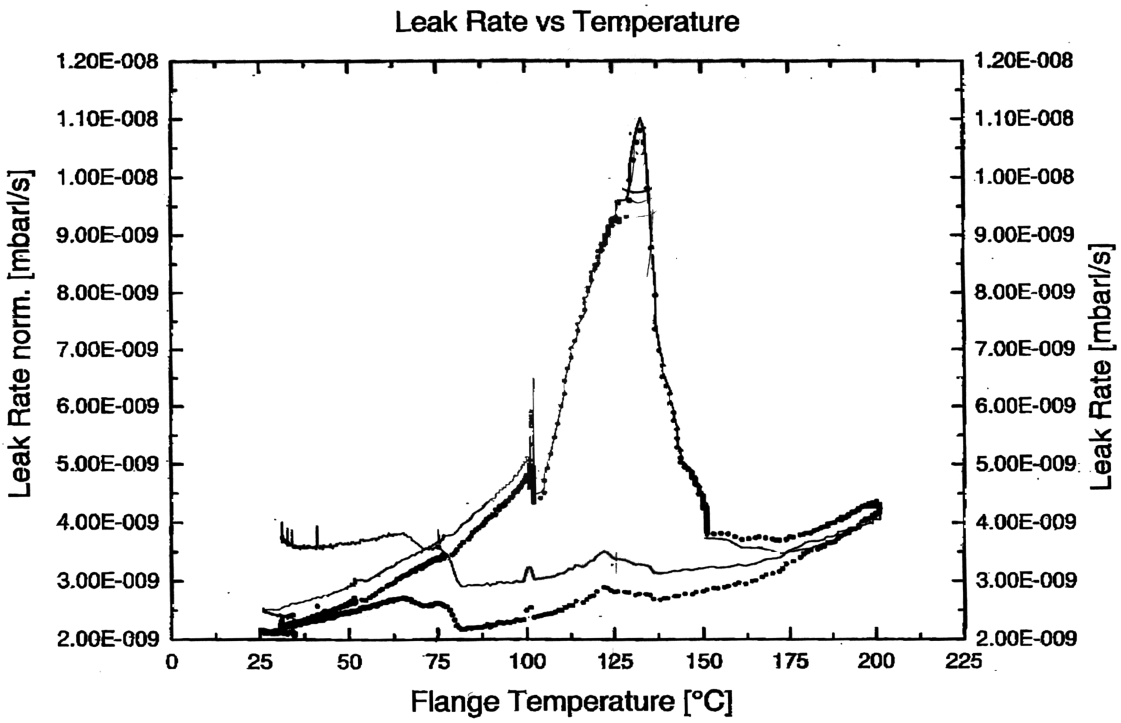

Figure 3.15 Internal data provided by Rob Church on leak rate vs baking temperature for one of their epoxy-sealed cells. There is a clear peak in leak rate of around half an order of magnitude around 125 °C (expected $= T_g$) and that the peak covers $T_g \pm 25$ °C.

to reach the baseline readings of the ion pumps ($< 10^{-10}$ mbar).

Optics 3.2

In all cold atoms experiments lasers and other optical elements are essential. Fortunately in our case we only need a relatively modest laser setup to achieve our goals. We require around 600 mW of 780 nm light (resonant with the Rb D2 line) to perform laser cooling as well as imaging of the atoms and around 40 W of 1550 nm light to perform evaporative cooling and to form an optical lattice.

As the 780 nm light is used to address atomic transitions the lasers producing this light need to have high specifications regarding linewidth and stability. Extended Cavity Diode Lasers (ECDLs) are a common laser architecture and can be used to produce several 10s of mW with linewidths < 1 MHz which are adequate for this work. These can be used as a seeding beam for a fibre- or tapered-amplifier, bringing

the output to several W.

Fortunately, high power lasers are readily available at 1550 nm as this is the standard wavelength used in telecom fibres due to silica's low absorption around this wavelength. Commercial units typically comprise a small seed laser which pumps a length of doped fibre, providing amplification, which is coupled to the unit's output. Such devices can provide 10s W of power at a relatively low cost per W compared to other wavelengths.

Locking and distribution 3.2.1

The laser beams which are used to address atomic transitions need to be stabilised in frequency against thermal drifts and mechanical shocks in order to stay resonant with the atoms. This is a standard procedure and there are many schemes to realise this, each providing different benefits. One of the most flexible is that of using a "master" laser, locked to an arbitrary transition, from which several "slave" lasers can be locked at fixed frequency offsets from the master laser frequency. The master laser does not need to be used for any other purpose and is often used merely as a reference oscillator for the optical system. As there are no constraints on which transition the master laser should be locked to the strongest available transition within a spectrum is usually chosen.

Saturated absorption spectroscopy

In previous sections we have assumed that we have a laser precisely tuned to an atomic resonance but have ignored the technical details as to how this is achieved. Section 2.2.1 alluded to difficulties in the application of traditional (linear) spectroscopic techniques due to the large Doppler broadening of the atomic absorption lines and thus we must look to nonlinear spectroscopy to resolve these.

One such technique commonly employed in atomic physics experiments is called saturated absorption spectroscopy, where an intense "pump" laser beam is used to

excite many atoms on a given transition (saturating it) whilst simultaneously using a weak, counter-propagating "probe" beam to measure the modified absorption due to these saturated transitions. These beams are generally derived from a single laser, allowing them to be synchronously swept in frequency in order to perform spectroscopy. As the pump and probe beams are counter-propagating they will address atoms with opposite velocity classes due to the Doppler shift required to bring each beam into resonance, however when the laser is tuned to a resonance both beams address the same near-zero velocity class. One can think of these atoms as having a very low effective temperature along the axis of the beams due to their low velocity component along this axis, which leads to a negligible Doppler broadening. Using relatively low laser powers prevents excessive power broadening and allows for the resolution of features on the order of the natural linewidth when the laser linewidth is sufficiently narrow (for ^{87}Rb, $<$6 MHz which is readily achievable with diode lasers).

Taking the first differential of a saturated absorption spectroscopy signal gives a so-called "error signal" which changes sign sharply on either side of the resonance and is zero on resonance. The first differential of the spectroscopy signal can be approximated electronically in real-time using a mixer to demodulate the spectroscopy signal with the sweep frequency. Due to the change of sign of the error signal when the laser frequency crosses a resonance the error signal can be used as the input to a feedback loop which controls the laser frequency (via a combination of altering the current supplied to the laser diode and the voltage supplied to piezo elements which control the laser cavity length) in order to "lock" the laser's frequency to an atomic resonance. Such "laser locks" are required to remove long term drifts of the laser frequency due to temperature fluctuations as well as short term drifts due to mechanical vibration or power supply noise. A feedback loop with sufficiently high bandwidth and error signal gradient can even reduce the linewidth of the laser when locked compared to its free-running state.

To see how we can extract the first derivative of our spectroscopy signal electronically we first define a function which takes the frequency of a laser beam, f, and returns a photodiode voltage, $V_{PD}(f)$, corresponding to the result of some spectroscopy. Secondly, we can write our laser's frequency as it is swept as a function of time as $f(t) = f_0 + \Delta f \sin(2\pi f_m t)$, where f_0 is the frequency of the laser without modulation, Δf is the modulation depth and f_m is the modulation frequency. We do not need to know the form of the function $V_{PD}(f)$ in order to extract information about its derivatives as we can approximate a function based on its derivatives using the McLaurin series,

$$F(x) = \sum_{n=0}^{\infty} \frac{f^{(n)}(x)}{n!} x^n, \tag{3.28}$$

where $F(x)$ provides a representation of our function $f(x)$ in terms of an infinite sum of the n^{th} derivatives of the function ($f^{(n)}(x)$) evaluated at some x. Taking our function $V_{PD}(f)$ and evaluating at $f(t) = f_0 + \Delta f \sin(2\pi f_m t)$, we have

$$\begin{aligned}
V_{out}(t) &= \sum_{n=0}^{\infty} \frac{V_{PD}^{(n)}(f(t))}{n!} f(t)^n \\
&= \frac{V_{PD}^{(0)}(f(t))}{0!} f(t)^0 + \frac{V_{PD}^{(1)}(f(t))}{1!} f(t)^1 + \dots \\
&= V_{PD}(f(t)) + V'_{PD}(f(t)) f(t) + \dots \\
&= V_{PD}(f(t)) + V'_{PD}(f(t)) \left(f_0 + \Delta f \sin(2\pi f_m t)\right) + \dots \\
&= V_{PD}(f(t)) + f_0 V'_{PD}(f(t)) + \Delta f \sin(2\pi f_m t) V'_{PD}(f(t)) + \dots,
\end{aligned}$$

which contains a term composed of the product of the laser modulation and the first derivative of the spectroscopy signal. We can now see that by passing our spectroscopy signal through a frequency mixer with the laser modulation signal we can shift terms proportional to the modulation frequency to DC (i.e. $f_m - f_m = 0$). Low-pass filtering the output signal of the mixer will remove any components remaining at harmonics of the modulation frequency resulting from higher order terms. This filtered output signal is now centred on DC and is proportional to

the first derivative of the signal seen at the photodiode versus frequency, $V_{PD}(f)$, which can be used to lock the laser's frequency to an atomic resonance. In order to observe the absorption spectrum we initially modulate the laser with a $f_m = 20$ Hz sine wave (higher modulation frequencies were found to add a number of spurious effects due to the response of the piezo element within our VECSEL). Once we have centred the laser frequency on the desired absorption line we reduce the amplitude of the modulation before switching the laser's modulation to the output of our PID controller.

In addition to the true absorption lines additional features called crossover resonances appear midway between two transitions connected by a ground or excited state. For the case of two excited states connected by a common ground state the pump beam depletes the ground state of atoms and populates one of the excited states leaving fewer atoms to absorb the probe beam which is resonant with the transition to the other excited state, as shown in figure 3.16(a). This situation causes a "positive" crossover resonance where we see an increased transmission at the midpoint of the two transitions. Conversely, a "negative" crossover resonance appears in a system with two ground states and one excited state, where the pump beam transfers the population of one ground state into the excited state which can decay into the other ground state causing a decrease in transmission at the midpoint of the two transitions.

Offset locking

With the master laser locked to an atomic transition we can divert a fraction of this beam's power to be overlapped with that of a slave laser on a photodiode. This forms a beat note between the two lasers with a characteristic frequency on the order of the ground state hyperfine splitting, which in the case of the rubidium is around 6.8 GHz. This frequency lies in the radio band of the electromagnetic spectrum and can be measured using off-the-shelf RF electronics components and interconnects.

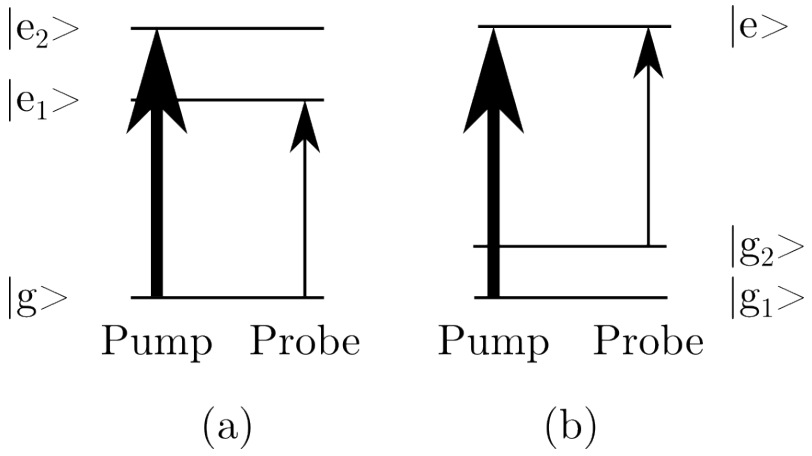

Figure 3.16 A schematic representation of the pump and probe beams interacting with two different states due to being resonant with two equal and opposite velocity classes. When the frequency of the pump are probe beams is set to the midpoint between the two transitions, this causes the pump beam to be resonant with one transition and the probe beam to be resonant with the other. The level structure in (a) gives rise to a "positive" crossover, where we see an increase in transmission and the level structure in (b) gives rise to a "negative" crossover where we see a decrease in transmission.

The detected frequency can be directly compared to an RF oscillator to provide an error signal with which to offset lock the slave laser or it can be mixed down to a lower frequency to be compatible with standard (i.e. non-RF) electronics.

In our case, we used a side-of-filter technique[53] to offset lock our lasers. This begins with converting the direct RF beat note to a lower frequency by mixing this with a voltage-controllable RF oscillator. This lower frequency is then split into two identical paths containing several stages of low-pass filters, with one path containing an additional high-pass filter at the start of the signal chain. The addition of carefully oriented diodes in each path allow for a DC signal to be produced after the paths are recombined which provides an error signal for locking. The benefit of this technique over others is that it produces a near-perfect error signal in that the signal sharply changes sign on either side of the locking point (where the signal is zero) but does not then decrease towards zero like typical error signals. This allows the lock to recover from very large excursions from the lock position, making this lock very robust. The offset amount can be changed by tuning the variable

RF oscillator at the start of the signal chain thus this setup can be duplicated for multiple lasers.

Fibre distribution

With the master laser locked to an atomic transition via spectroscopy of Rb vapour a slave laser is locked at a fixed frequency offset from this using the method described above to provide the requisite laser beams to perform laser cooling and recycling of atoms excited to "dark" states. These two initial beams ("cooling" and "repumping" respectively) need to be divided into several paths and mixed with various ratios in order to realise the final beams we require to produce our 2D and 3D MOT beams, pushing beam and imaging beam. To do this we opted to use a fibre-based system as opposed to free space in the hope of achieving a system which is stable against polarisation fluctuations which are often present in free space systems due to issues with air flow and temperature fluctuations along beam paths. Splitting and overlapping of beams with varying power ratios can be achieved by means of splices between two fibres running parallel to one another. The outer cladding of the fibres are stripped back and the evanescent coupling of light from one fibre to another can be tuned depending on how closely the two fibres are positioned. A variety of optical components can be emulated using such a coupler, as illustrated in figure 3.17, meaning that complex optical setups can be constructed using these couplers which will not require alignment. From these basic building blocks I designed two evanescent fibre coupler networks which would provide the majority of our optical setup. One provides much of the mixing and splitting of the input beams and the other is specifically for the 3D MOT, where we wish to have a precise control over the power balance of each of the beams. There are unavoidable power losses at each evanescent coupler on the order of 5 %, however these have been accounted for in the fibre network design and as such we are able to achieve sufficient power at all of the outputs to perform our experiment. A schematic of the optical setup

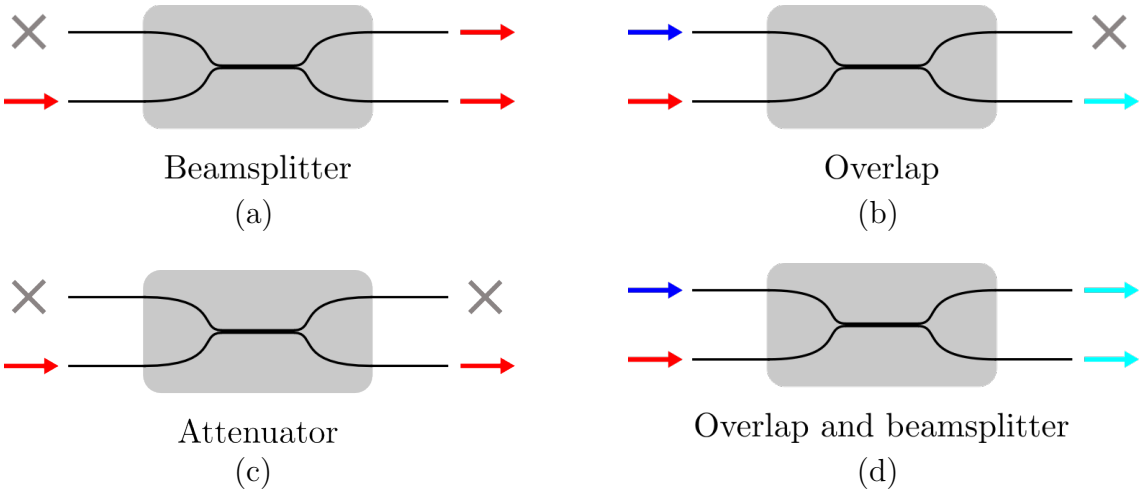

Figure 3.17 Different configurations of inputs and outputs for a 2-in/2-out evanescent coupler which form several common building blocks of optical setups.

implemented by these evanescent coupler networks is shown in figure 3.18.

Dipole trap 3.2.2

We had opted to use a single pass dipole trap in order to perform evaporative cooling on the atoms captured by the 3D MOT. A 10 W 1550 nm laser was chosen for this and is focused down to a waist of 42 μm using a lens with 50 mm focal length. As this beam is fibre-coupled we expect to lose a percentage of this light, perhaps limiting the achievable power at the fibre output to 7 W. The profile of the beam leads to strong trapping in the focal plane and weak trapping along the axis of the beam. It is anticipated that the weak trapping along the beam may be detrimental to achieving BEC and thus I have allowed for room to recycle the beam exiting the cell into a crossed-beam arrangement, deepening the trap along the weak axis by providing a second focal plane at an angle to the original focal plane.

Optical lattice 3.2.3

In addition to affording a greater range of experiments we can perform, such as those on magnetically frustrated systems[21–28], an optical lattice is to be used in our

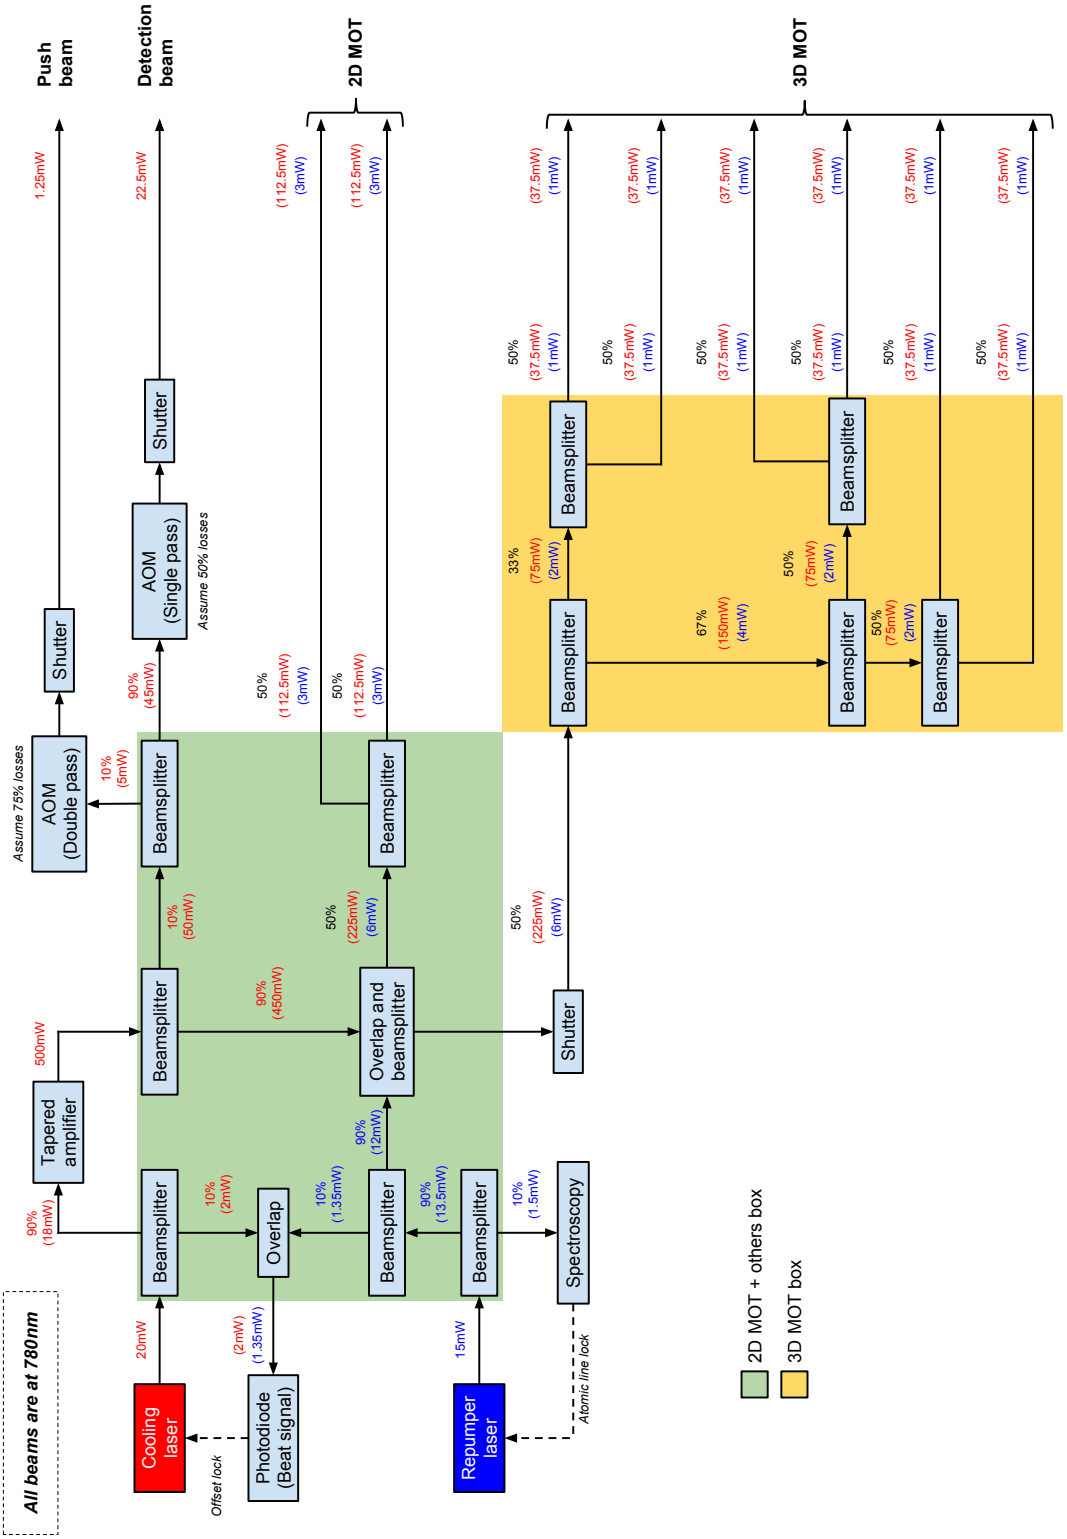

Figure 3.18 A schematic diagram of the optical setup for our 780 nm beams. Much of the mixing and overlapping of beams can be performed within the evanescent coupler network, meaning that the majority of the setup does not require alignment.

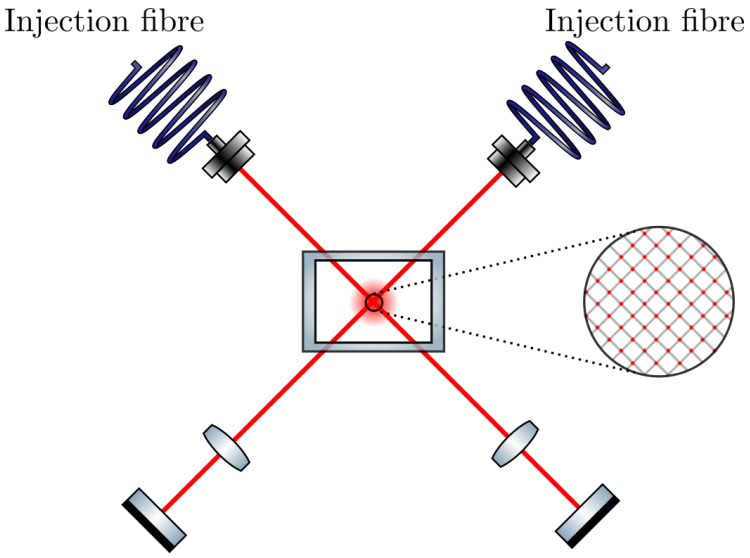

Figure 3.19 The optical configuration for our two "normal" lattice beams. Each beam comes out of a fibre and focuses onto the atom cloud (due to focusing optics on the fibre output) and is then collimated using a lens, reflected by a mirror and refocused onto the atoms using the same lens. On the right, a detailed view shows the peaks of the standing waves created by the two laser beams (*red*).

work to enhance the strength of dipolar interactions through long-range collective effects[41]. Our 3D optical lattice is different from typical beam arrangements in that it is not constructed from 3 orthogonal beam pairs. Instead there are two orthogonal retroreflected beam pairs, shown in figure 3.19, overlapped with a further beam which is directed such that it intersects itself at an angle, shown in figure 3.20. The result of this intersection produces a standing wave pattern in the overlap region, with a period depending on the angle at which it intersects itself. The benefit of this arrangement is that all beams travel through the same face of the glass cell, leading to a reduction in the complexity (and thus cost) and increase in performance of the optical coating on the cell.

Imaging 3.2.4

With such tight constraints on space around the glass cell (in order to fit inside our passive magnetic shielding) as well as the allowed magnetic field levels in this region (on the order of 10s pT) we are precluded from using any electronics within

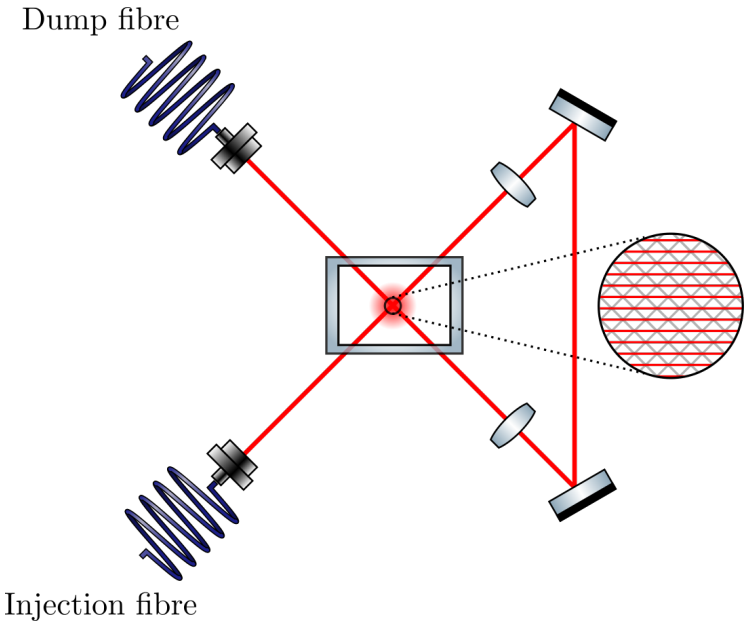

Figure 3.20 The optical configuration for our two "split" lattice beams. The beam comes out of the fibre and focuses onto the atom cloud (due to focusing optics on the fibre output) and is then collimated and steered back to the incidence face of the glass cell using two mirrors, focused onto the atom cloud again using a lens and is finally coupled into a fibre whose output is incident on a beam dump outside of our magnetic shield. On the right, a detailed view shows the peaks of the standing waves (*red*) created by the wavefronts of the two components of the laser beam (*grey*).

the shield. This can be understood by calculating the field at a distance from a current-carrying wire, where a current of 1 μA at a distance of 5 cm produces a flux density of 4 pT. Whilst it may be possible to keep any wires and camera >5 cm away from the BEC, sufficiently small and low-power (i.e. low current) cameras are not available with which to image it without producing deleterious stray magnetic fields. Another option would be to include holes within the passive shielding in order to allow an image of the BEC to pass through and onto a camera, however this is highly undesirable as will allow ambient fields to penetrate into the shield, as discussed in section 3.3. For these reasons we have opted to image the BEC onto a coherent fibre bundle which is able to transmit the image outside the shield onto a camera. This allows us to only have passive components inside our shielding and enables us to use any camera we wish. With the advancements in medical imaging, coherent fibre bundles are now readily available off-the-shelf with up to 18,000 fibres in a 1.65 mm diameter bundle, with individual fibres as small as 6 μm in diameter[15]. This is comparable to the pixel pitch of many cameras and thus we should not be adversely affected in terms of resolution by using this method.

Non-magnetic optics 3.2.5

In order to position optics around our glass cell we require some sort of optical breadboard in place. However, as we will surround this area with passive shielding it cannot be mounted to the underlying optical table as this would conflict with the shielding. In addition, we require optics to be placed in several planes around the glass cell to implement our 3D MOT and optical lattice. To overcome these issues I designed a configuration of optical breadboards, shown in figure 3.21, which provide the planes we need, fit within the innermost passive shield, and mount to an arm bolted to the vacuum chamber. All of these parts are made from aluminium to provide sufficient mechanical strength whilst being non-magnetic. We have chosen to

[15]Available as leached fibre bundles from SCHOTT.

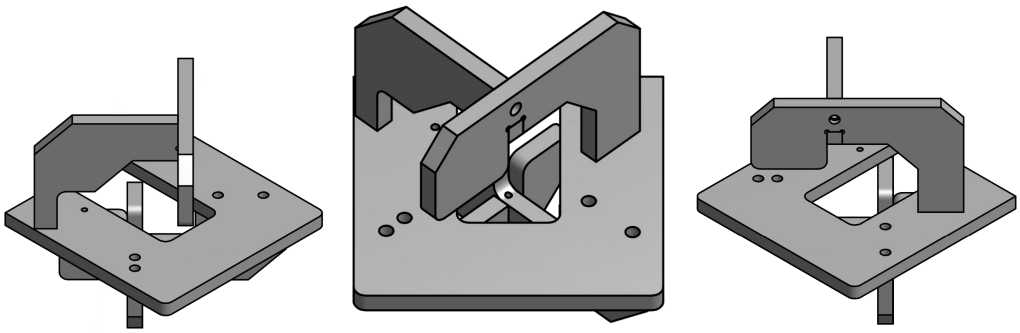

(a) (b) (c)

Figure 3.21 Optical breadboards designed to fit around our glass cell and within our passive shielding. From the views from (a) the left and (c) right (less so in the (b) centre view) it can be seen that the upper and lower angled planes are offset from one another. This is to allow easier mounting of these planes to the horizontal plane. The offsets of each plane are calculated such that the optic-centre height of the mounts is in the same plane as the atoms such that beams in all planes will overlap on the position of the atoms.

use high strength epoxy to hold our optics in place to avoid the need for additional clamps due to our tight space constraints. To enable this, the positions of these breadboards relative to the glass cell have been carefully chosen to be equal to the height of the optic-centre of the mounts we have chosen to use, such that no height adjustment is necessary.

Similar to our constraints on imaging, we also must take care in our choice of optics mounts in terms of size and stray magnetic fields. Standard optical mounts typically have steel ball-bearings and springs in order to provide kinematic adjustment, which could add stray magnetic fields to our experiment from their residual magnetisation. Further, standard 1/2" optics mounts are sufficiently bulky to cause packing density problems within the tight confines of our shield. For optics requiring kinematic adjustment we opted to use low-profile flexure mounts[16], however we looked towards 3D printing to produce bespoke, low-profile, non-magnetic optics mounts for the non-kinematic optics around the glass cell.

3D printing is a rapidly advancing technology and has become sufficiently cheap within the past few years to enable hobbyists and enthusiasts to buy or have access

[16]Purchased from Siskiyou.

to 3D printers. For our purposes, 3D printing provides a means to produce bespoke parts from an inherently non-magnetic material with a sufficiently high resolution to make the low-profile mounts we need. One part which was unable to be sourced in a non-magnetic material or sufficiently small was a rotational mount, used to hold polarisers and waveplates. Part of the difficulty of this is that these optics are typically not available in smaller sizes than 1/2". Fortunately, we were able to have large versions of these optics (ca. 2" diameter) precision cut into 5 mm squares[17]. These squares were able to be held in a small wheel (see figure 3.22) using a small amount of epoxy applied with a toothpick. This retaining wheel rests in a larger mount with a hole, allowing for light to pass through the optic. The wheel mates with the mount with a tight tolerance but is able to rotate freely and once adjusted can be secured with a captive nylon nut and screw. The height of the centre of the optic was chosen to be equal to that of the flexure mounts we are using so that no height adjustment is required for use. An additional benefit of 3D printing is the availability of multiple colours of plastic, enabling us to colour-code our optics as to their function (i.e. polarisers or waveplates). These parts and their assembly can be seen in figure 3.22.

Passive magnetic shielding 3.3

The ultimate goal of our experiment is to study dipole-dipole interactions between Rb atoms within a spinor BEC. The energy of these interactions is on the order of $E_{dd}/k_B \approx 0.1$ nK and thus we must ensure that ambient noise sources are below this level. Particularly detrimental to these studies are stray magnetic fields which need to be reduced to a level of around 10 pT in order that their resulting energy shifts are negligible in our case. This is no easy task as Earth's magnetic field is around 50 μT and fluctuates to at a level of several hundred nT over timescales of minutes and at a level of several μT over hours. In addition most lab equipment

[17]Cutting performed by Fraunhofer Institute.

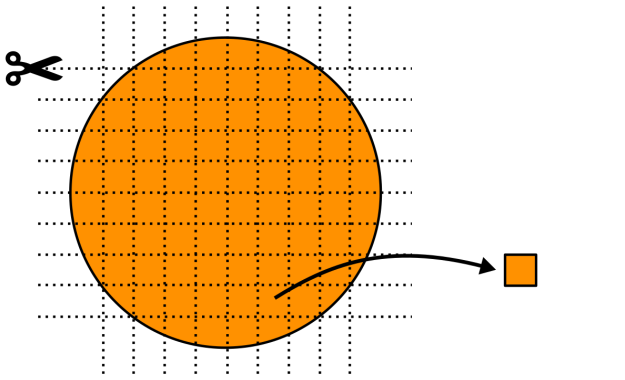

2" diameter standard optic 5 mm square optic

Empty wheel

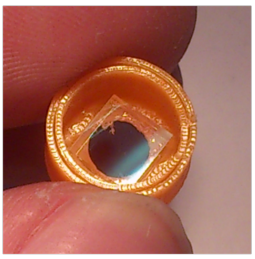

5 mm square optic mounted in wheel

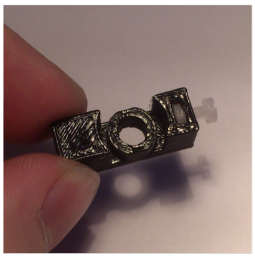

Empty base

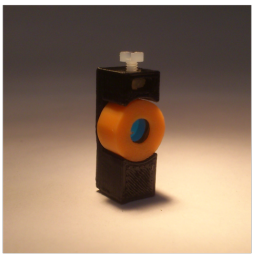

Wheel-mounted optic mounted in base

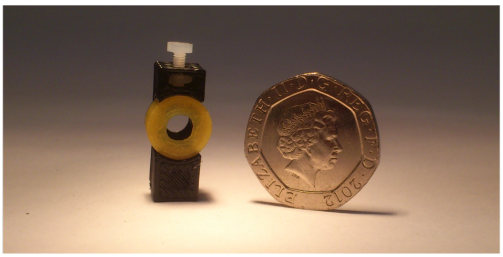

20 p piece for scale

Figure 3.22 A schematic showing how I assembled the 3D printed miniature optics. First, a standard size circular optical element (waveplate or polariser) is precision cut into 5 mm square sections, which are held in a 3D printed "wheel" with some small spots of epoxy. This "wheel" is then held in a base by a friction fit. A hole in the mounting feature in the base allows for light to pass through the base and in turn through the optic. The wheel is held securely by friction and is just able to rotate. A small nylon nut and screw are held captive within the base which grip the wheel in order to lock its rotation once aligned.

and building infrastructure contributes further noise e.g. from mains conduits, large current pulses, moving vehicles, elevators, trains, and subway cars. Many of these sources are beyond reasonable control and can pose serious problems for precision measurements.

Magnetic fields can be reduced within a given volume by surrounding it with a shell made from a material with high relative permeability, μ_r. Shielding magnetic fields in this way is termed passive magnetic shielding, in contrast with active magnetic shielding where a sensed field is produced in anti-phase using electromagnets driven by a negative feedback system. Passive magnetic shielding is one of the only ways to reduce DC magnetic fields to the levels which we require and as such I have designed a state-of-the-art passive shielding system for our experimental chamber. Magnetic fields are reduced by these shields due to flux shunting of low frequency fields and to a lesser extent by eddy current shielding via the skin effect of high frequency fields. Flux shunting is a consequence of the magnetic field, $\mathbf{H}$, and resulting flux density, $\mathbf{B}$, obeying boundary conditions across materials with different relative permeabilities, such as between the shield material and surrounding air, which we will discuss in detail in the following section. However, as the majority of our noise sources are AC or transient we have opted to also use an active magnetic field stabilisation system in conjunction with our passive magnetic shielding to further improve the noise immunity of the experiment, which will be discussed in detail in chapter 4.

Theory overview 3.3.1

A boundary condition is imposed on the component of the flux density perpendicular to a boundary and can be derived by considering the magnetic flux through a Gaussian volume (shown in figure 3.23) which extends across the boundary between two materials. Maxwell's second law states that $\nabla \cdot \mathbf{B} = 0$, physically meaning that for all flux lines entering a volume there must be an equal number which exit. In

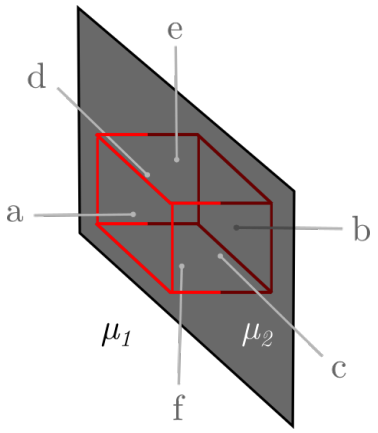

Figure 3.23 Illustration of a Gaussian volume extending across a boundary between two materials with different relative permeabilities, μ_1 and μ_2.

the limit that the extent of this volume across the boundary becomes vanishing the sole contribution to the net flux divergence then becomes that of the faces normal (a) and anti-normal (b) to the boundary as the area enclosed by the other faces (c, d, e, f) approach zero. Thus the flux lines entering the normal face must all exit the anti-normal face. This leads to the normal component of $\mathbf{B}$ being continuous across the boundary as the number of flux lines and area, A, on each side are equal. This can also be seen by applying the integral form of Maxwell's second law to the closed Gaussian volume, S:

$$\begin{aligned}
\nabla \cdot \mathbf{B} &= 0, \\
\therefore \oint_S \mathbf{B} \cdot \mathbf{dA} &= 0, \\
&= \left[\int_a + \int_b + \underbrace{\int_c + \int_d + \int_e + \int_f}_{=0} \right] \mathbf{B} \cdot \mathbf{dA} = 0, \\
&= \left[\int_a + \int_b \right] \mathbf{B} \cdot \mathbf{dA} = 0, \\
&= B_a A - B_b A = 0, \\
\therefore B_{\perp \text{in}} &= B_{\perp \text{out}}.
\end{aligned} \tag{3.29}$$

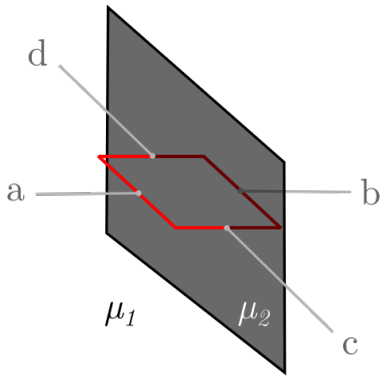

Figure 3.24 Illustration of an Amperian loop extending across a boundary between two materials with different relative permeabilities, μ_1 and μ_2.

A second, separate boundary condition applies to the component of magnetic field orthogonal to a boundary normal by considering an Amperian loop (shown in figure 3.24) which spans the boundary between two materials. Using Maxwell's fourth law, $\nabla \times \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{E}}{\partial t}$, and assuming that we are dealing with magnetostatics (i.e. $\frac{\partial \mathbf{E}}{\partial t} \to 0$) and that there are no surface currents (i.e. $\mathbf{J} \to 0$) then we have that $\nabla \times \mathbf{H} = 0$. A physical understanding of this effect is harder to visualise but arises from the fact that magnetic fields are conservative potentials and any force exerted on a particle travelling through a field along a given path is equal and opposite to that exerted upon reversing the direction of the path. In the limit that the length of the loop becomes vanishing in the normal direction of the boundary this causes the contribution to the net curl of the magnetic field due to the loop edges perpendicular to the boundary (c and d) to tend to zero. In a similar fashion as before, this leads to the result that the component of the magnetic field orthogonal to the boundary normal, $\mathbf{H}_{\parallel}$ (i.e. *parallel* to the boundary as opposed to *orthogonal* to the boundary), is continuous across the boundary. This can also be seen by applying the integral

form of Maxwell's fourth law to the closed Amperian loop (C) we have described:

$$\begin{aligned}
\nabla \times \mathbf{H} &= 0 \\
\therefore \oint_C \mathbf{H} \times \mathbf{dl} &= 0 \\
&= \left[\int_a + \int_b + \underbrace{\int_c + \int_d}_{=0} \right] \mathbf{H} \cdot \mathbf{dl} = 0 \\
&= \left[\int_a + \int_b \right] \mathbf{H} \cdot \mathbf{dl} = 0 \\
&= H_a l - H_b l = 0 \\
\therefore H_{\parallel \text{in}} &= H_{\parallel \text{out}}
\end{aligned} \tag{3.30}$$

From the two boundary conditions derived in the preceeding paragraphs we can now see how high relative permeability materials are able to shield magnetic fields. As only the parallel component of the *field* is continuous across a boundary this causes the parallel component of the *flux density*, $\mathbf{B}_\parallel = \mu \mathbf{H}_\parallel$, to be discontinuous by an amount equal to the ratio of the relative permeabilities of the two materials ($\Delta\mu = \frac{\mu_2}{\mu_1}$). In the case of a boundary between air and μ-metal ($\Delta\mu \approx 10^4$) this causes a large deviation of the direction of the flux density vector when it passes from air into μ-metal and leads to the majority of incident flux being channeled through the shielding material rather than penetrating it, as illustrated in figure 3.25. The figure of merit used to measure field reduction is the shielding factor, defined here as the ratio between the flux density outside the shield to that inside, $S = \frac{B_{\text{ext}}}{B_{\text{int}}}$. As the performance of magnetic shields varies with frequency it is more useful to quote the shielding factor separately for DC and AC fields, S_{DC} and S_{AC} respectively. Shielding factors as high as $S_{\text{DC}} = 10^8$ and $S_{\text{AC}} = 10^{14}$ have been demonstrated for carefully designed and constructed shields such as those used on Gravity Probe B[54].

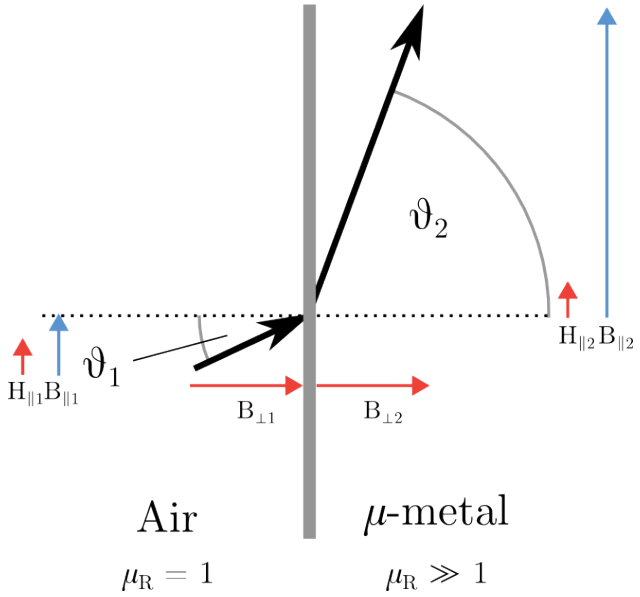

Figure 3.25 A schematic diagram of the magnetic field components of a field line **B** across a boundary. The perpendicular component of **B** is continuous across the boundary, but only the parallel component of **H** is continuous across the boundary. These continuous quantities are highlighted in red in the diagram. As only **H** is continuous parallel to the boundary, this causes the parallel component of **B** to be discontinuous by an factor of $\Delta\mu = \mu_{R2}/\mu_{R1}$, where μ_{R1} and μ_{R2} are the relative permeabilities of the materials on either side of the boundary respectively.

Previous work

Analytical solutions for the shielding factor of nested shields have been sought after since the first formal calculations in 1894 by Rücker who studied the case of concentric spherical shells[55]. Later investigations included those of nested spherical shells[56], nested cylindrical shells of infinite length[57] (in order to exclude end effects, practical shields may approach this limit as the length to diameter ratio becomes very large), open-ended nested finite length cylindrical shields[58], closed nested cylindrical shells of finite length[59], and nested cylindrical shells with holes along their axis[60] which broadly describes most shields in use today.

Using the recursive equations derived by Sterne[57] for nested cylindrical shields of infinite length we can estimate the shielding factor, S, of a set of nested shields:

$$S = \frac{1}{2}\left(u_{n+1} + v_{n+1}\right)$$

$$u_{n+1} = \alpha_n u_n + \beta_n v_n$$
$$v_{n+1} = \gamma_n u_n + \delta_n v_n$$

$$\alpha_n = 1 - \epsilon_n - \epsilon_{n,n+1} + \left(\mu_n + 1\right)\epsilon_n \epsilon_{n,n+1}$$
$$\beta_n = \epsilon_{n,n+1}\left(1 - \epsilon_n\right) + \left(\frac{\epsilon_n}{\mu_n}\right)\left(1 - \epsilon_{n,n+1}\right)$$
$$\gamma_n = \mu_n \epsilon_n + \epsilon_{n,n+1} - \left(\mu_n + 1\right)\epsilon_n \epsilon_{n,n+1}$$
$$\delta_n = 1 - \epsilon_n - \epsilon_{n,n+1} + \left(\frac{\mu_n + 1}{\mu_n}\right)\epsilon_n \epsilon_{n,n+1}$$

$$\epsilon_n = \left(\frac{RO_n - RI_n}{RO_n}\right) - \frac{1}{2}\left(\frac{RO_n - RI_n}{RO_n}\right)^2 \tag{3.31}$$
$$\epsilon_{n,n+1} = \left(\frac{RI_{n+1} - RO_n}{RI_{n+1}}\right) - \frac{1}{2}\left(\frac{RI_{n+1} - RO_n}{RI_{n+1}}\right)^2 \tag{3.32}$$

where equations (3.31) and (3.32) can be re-cast as

$$\epsilon_n = a - \frac{1}{2}a^2$$
$$\epsilon_{n,n+1} = b - \frac{1}{2}b^2$$

$$a = \frac{d_n}{RO_n}$$
$$b = \frac{x_n}{RI_{n+1}}$$

where d_n is the thickness of the n^{th} shield and x_n is the inter-layer spacing from shield n to shield $n+1$ (x_n is defined as zero for the outermost shield). These equations were evaluated for various radii of innermost shield for constant spacings between the shields, seen in figure 3.26(a), and for a geometric progression of the shield spacings, seen in figure 3.26(b). It is clear that there are optimal values for both the spacing and geometric progression factor. Both graphs indicate that smaller shields provide higher shielding factors, as we expect, however we are limited by space constraints to a minimum innermost shield radius of 10 cm. A slice through each of these plots for $r = 10$ cm can be seen in figure 3.27 which show that geometric progression appears to yield a higher shielding factor (by around a factor of 2). In light of this, we have chosen the sizes of our shields around this optimal value of the geometric progression factor based on an innermost shield radius of 10 cm.

It is commonly cited that spherical shields provide the highest shielding factor of common shapes[61–64], being infinte for a closed shell[65]. However, due to the difficulty of manufacturing spherical shells cylindrical shells are more commonly used, where the shielding factor improves as the length to diameter ratio approaches unity (i.e. as in the case of a sphere)[59,61]. Additionally it has been shown that the resultant shielding factor scales inversely with the contained volume, meaning that smaller shielded volumes give larger shielding factors[57,58,62]. This fact that has been exploited to construct highly effective shields for chip scale atomic magne-

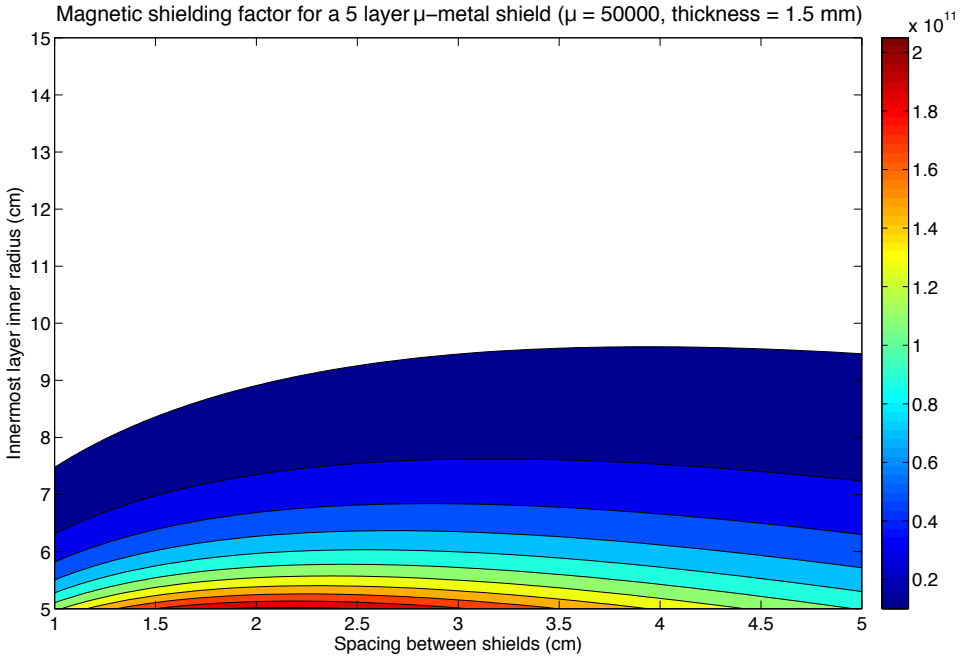

(a)

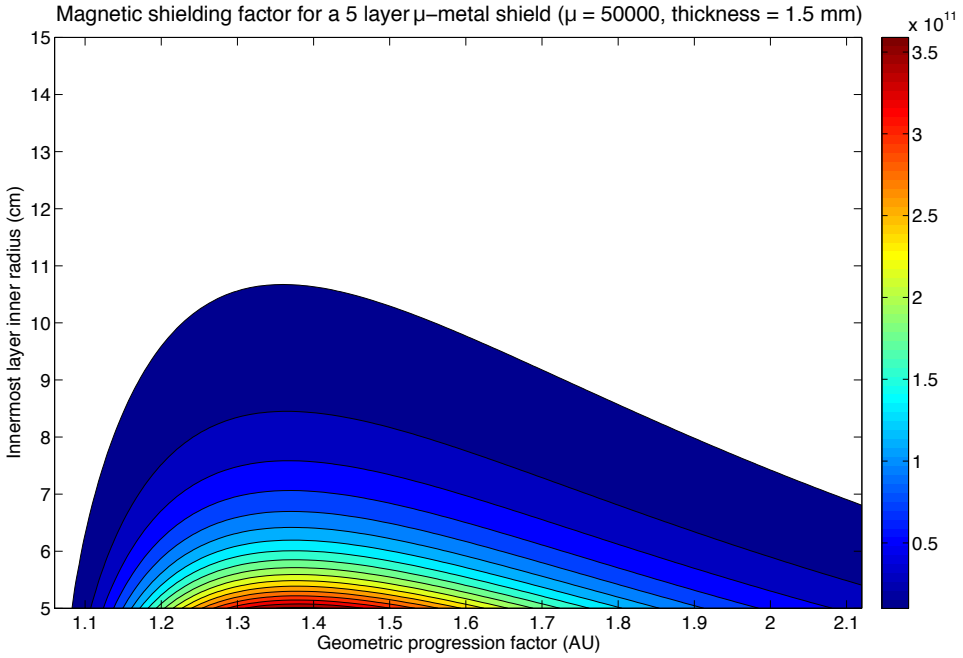

(b)

Figure 3.26 3D plots of the shielding factor, S, achieved with a 5 layer μ-metal shield made from 1.5 mm thick $\mu_R = 50,000$ material vs the inner radius of the innermost shield for (a) a constant spacing, d, between each of the shields and (b) spacings based on a geometric progression, with a co-efficient a. The colour scale for these graphs is terminated at 1×10^{10}, meaning that any white regions of the graph have $S < 1 \times 10^{10}$.

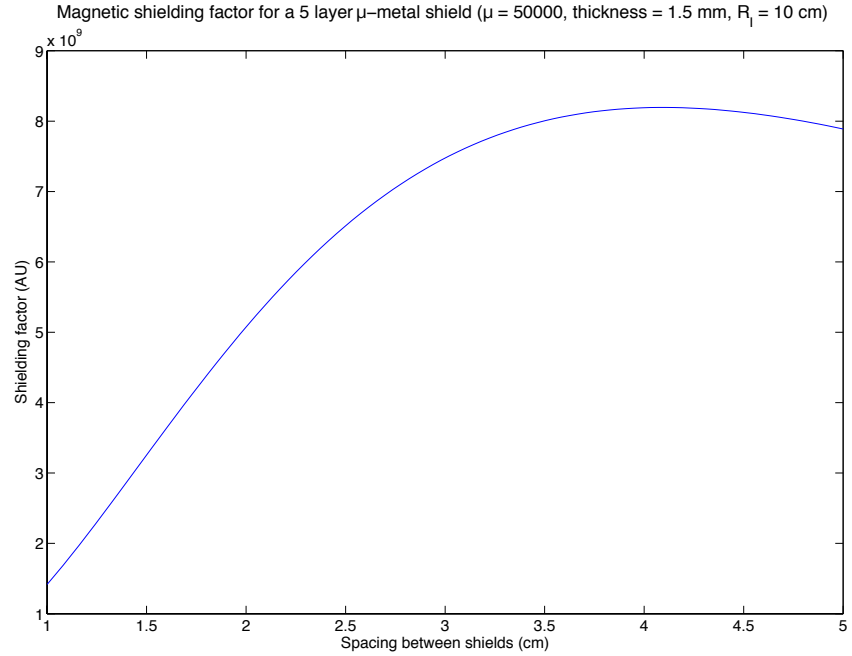

(a)

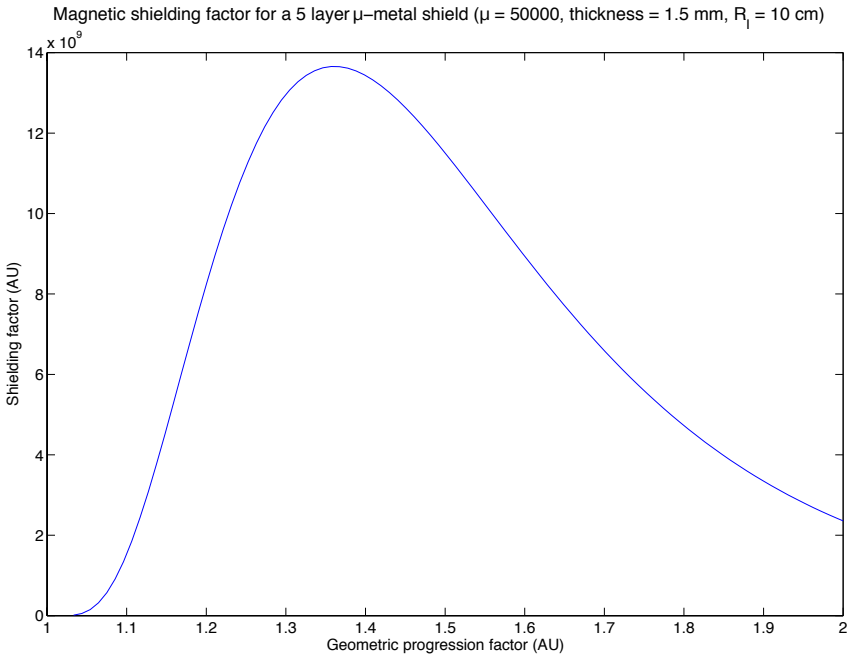

(b)

Figure 3.27 2D plots of the shielding factor achieved with a 5 layer μ-metal shield made from 1.5 mm thick $\mu_R = 50,000$ material with an inner shield inner radius of 10 cm for (a) a constant spacing between each of the shields and (b) spacings based on a geometric progression. The maximum shielding factor obtained is 1.4×10^{10} at $a = 1.357$ for (a) and 8.2×10^9 at $d = 4.05$ cm for (b).

tometers[66].

Whilst the case of open-ended cylindrical shields pertains to many practical shields, the open ends of the shield allow field lines to penetrate along the axis of the shield reducing the axial shielding factor. Such field penetration can be minimised by closing the ends of the cylinder, allowing some small holes in these caps for experimental access. Field lines can still permeate into the shield through this hole but decay exponentially with a power which depends on the relative orientation of the external field and the axis of the hole, ranging from 1.0–3.5[63]. The distance which the permeating field extends into the shield is proportional to the diameter of the hole and is quoted in various places as being somewhere between 2–5 times the hole diameter. For small holes, the reduction in shielding factor along the axis is quite small — on the order of 10% compared to a closed shield[67]. Field permeation into holes in the shield can be mitigated by welding short tubes of the same diameter of the hole to the hole, called tubulations, whose length is recommended to be 2–5 times the hole diameter[68]. It should be noted that there is little analytical treatment of the effect of tubulations in literature but rather this is largely an empirical result confirmed from several constructed shields[69]. Literature suggests that the axial shielding factor of nested shields with an axial hole increases towards the limit of closed shields for large axial separations of the shields[60]. Regardless of the mitigation strategies used, the axial shielding factor is generally the weak point of any shielding design and thus it is this which must be optimised in order to achieve satisfactory performance. By considering these various constraints we can begin to imagine the construction of a passive magnetic shield with a high axial (and thus, also radial) shielding factor — several concentric, well separated, closed cylindrical shields with tubulated axial holes.

Magnetic shields are commonly manufactured from materials such as μ-metal (a general class of metal alloy comprising iron, nickel, copper, chromium, manganese and molybdenum in varying amounts depending on the particular brand[70]), ceramic

ferrites (such as MnZn[71]), soft ferrites (such as iron) and good electrical conductors (such as copper and aluminium). One can also construct magnetic shielding from materials like Metglas, a high relative permeability alloy usually manufactured in very thin strips. Several strips can be wound around a supporting structure to create a passive magnetic shield[72]. The construction of large passive magnetic shields may require several welds along their length giving rise to a variation in shielding factor around these welds, but this can be mitigated by wrapping the welds with Metglas foil at the expense of field homogeneity within the shield[73]. The reason Metglas in particular disturbs the local field homogeneity is that when flux lines are channeled through the material they are forced into such a small cross-sectional area at the edges of the thin material, causing the flux density at the edges of the strips to be enormous. These large, localised field strengths necessarily give rise to local spatial field gradients. The unifying property of all but the latter is their high μ_r which can be anywhere between 10^2–10^6. Materials with high relative permeability provide shielding primarily through the flux shunting mechanism and materials with a high conductivity via induced eddy currents. In general, materials with high permeability are not good conductors and vice versa.

The use of good conductors for magnetic shielding utilises the phenomenon of the skin effect to reduce penetrating fields by exciting eddy currents within the material (readily done due to their high conductivity) which in turn generate impressed magnetic fields opposing the incident field, causing it to decay exponentially with distance in the conductor. The relevant length scale of the decay is the skin depth of the conductor, $\delta = \sqrt{\frac{2\rho}{\omega\mu}}$, at which the power will have dropped to $\frac{1}{e^2} \approx 13\%$ of its surface value. This phenomenon is strictly dynamical and does not apply at DC, as can be seen by the fact that $\delta \to \infty$ when $\omega \to 0$ or considering Faraday's law, $\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$, where $\frac{\partial \mathbf{B}}{\partial t} \to 0$ at DC and thus $\nabla \times \mathbf{E} \to 0$ giving that no opposing electric field is induced within the material at DC. These effects are frequency dependent and give rise to a higher shielding factor for higher frequencies, hence such

eddy current shielding using good conductors is typically used to reject RF fields and other EMI. Eddy current shielding is also present in μ-metal and the other shielding materials mentioned, albeit at a lower level due to the lower conductance, and thus these materials provide shielding of both AC and DC fields, with conductors only shielding AC. For typical high permeability shields eddy current shielding typically becomes dominant above 100 Hz[61,62].

Magnetic shielding can itself produce magnetic noise due to thermal fluctuations of charges within the shield[74,75] which proves to be a limiting factor in state of the art magnetometers[76]. As this magnetic noise is rooted in the conductivity of the shielding material, one can drastically reduce its effects by using a ceramic ferrite (which generally have very low conductivity) rather than a metal alloy. Only the innermost shield of a set of nested shields need be made of this ceramic ferrite material in order to reduce noise. This approach has been demonstrated to lower the magnetic noise floor by a factor of 25 compared to an equivalent mu-metal shield[71]. Thermal magnetic noise is typically on the order of 10s fT/$\sqrt{\text{Hz}}$ and thus these considerations are only necessary for the most sensitive of magnetometers and when making biomagnetic measurements.

Simulating shielding performance 3.3.2

As μ-metal shields are incredibly expensive, due to high material costs and the sophisticated tooling required for manufacture, it is common practice to perform computer simulations of the performance of a shield design before having it made. Whilst sufficiently detailed simulations can be complex and time-consuming to set up and run, these can provide results which agree very well with empirical measurements of the physically constucted shields and is thus highly recommended.

In order to calculate the magnetic flux density distribution over a geometry it is often easier to work in terms of the magnetic vector potential, $\mathbf{A}$, which is defined as $\mathbf{B} = \nabla \times \mathbf{A}$. Selecting the Coulomb gauge (i.e. $\nabla \cdot \mathbf{A} = 0$) and using magnetostatics

with no free currents (i.e. $\nabla \times \mathbf{H} = 0$) we see that

$$\begin{aligned} \nabla \times \mathbf{B} &= \nabla(\nabla \cdot \mathbf{A}) - \nabla^2 \mathbf{A}, \\ \therefore \nabla \times (\mu \mathbf{H}) &= -\nabla^2 \mathbf{A}, \\ \therefore \nabla^2 \mathbf{A} &= 0. \end{aligned} \tag{3.33}$$

One can predict the shielding factor for a given shield by solving the Laplace equation, $\nabla^2 \mathbf{A} = 0$, for the magnetic potential, $\mathbf{A}$, over a given volume whilst honouring the material boundary conditions discussed above. Such calculations, whilst simply formulated, give rise to very complex and recursive expressions for even simple shield geometries. The shielding performance of multiple shields exceeds the product of its constituent shields and thus using several nested shields is the most common method of reaching ultralow magnetic fields within the inner shield volume.

Shield design 3.3.3

Whilst the analytical approximations can aid in reaching a suitable region of parameter space for a given set of shields it is difficult to predict the performance for such complex geometry shields anayltically and thus numerical approximation techniques are used. A popular technique for solving computational problems where several free parameters are allowed to vary over space whilst honouring certain global and boundary conditions is finite element analysis. In finite element analysis the geometry of the problem is defined, a mesh applied over its surface and a trial solution is proposed. This trial solution is then compared to an equation which governs the entire system which is evaluated at each of the vertices of the mesh with respect to its neighbouring vertices. The trial solution is adjusted iteratively to better match the global and local boundary conditions in order to minimise a chosen convergence parameter which measures how well the trial solution matches a chosen target criteria. In order to simulate the shielding factor for a complex geometry shield one must

design the shield in a CAD program, apply a mesh over its surface and evaluate the magnetic potential over each mesh vertex whilst honouring the Laplace equation globally as well as the boundary conditions between materials with different relative permeabilities. The critical parameter determining the effectiveness of a passive magnetic shield is its relative permeability, which varies with field strength[58]. In a low external field the relative permeability can be reduced by as much as a factor of 10–100 compared to the permeability at high fields[77]. In order for simulations to more closely match experimental realisations this must be taken into account. In general, relative permeabilities of mu-metal alloys are quoted at a field strength of 40 G, quite far from zero field. There are many implementations of the exact solving configuration in terms of how the program proceeds towards a region of parameter space which minimises the convergence parameter, but for our purposes these are uninteresting to consider in detail.

Based on the considerations outlined above, we have designed a passive shielding system to be incorporated into our experiment. Broadly, the comprises 5 nested cylindrical shells which are closed at one end with tubulated holes along their axis on the other. These access holes allow for our glass cell and optical fibres to enter into the shielding assembly as well as images to be transferred out of the shielding assembly for processing. CAD drawings of the proposed shield design can be seen in figure 3.28.

Simulations 3.3.4

Simulations of the shielding factor of the shield design presented in figure 3.28 were performed using COMSOL Multiphysics. Such simulations involve replicating the CAD model within the software, defining the problem's boundaries and materials involved, and finally creating a sufficiently fine mesh over the problem geometry for finite element analysis upon. Due to the small details involved (i.e. mm thickness shields and edge features) over large scales (i.e. shields with diameters of several

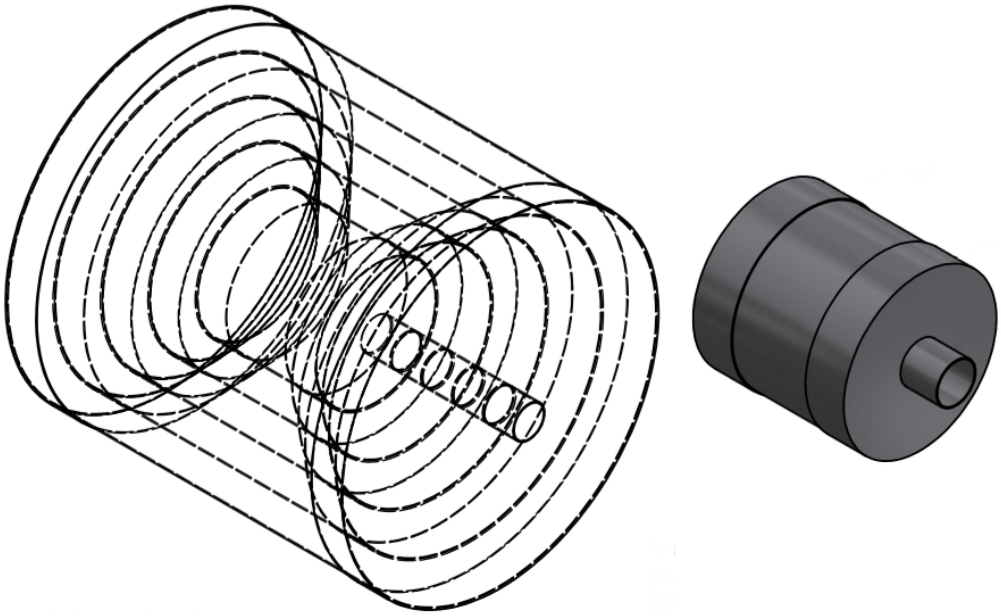

Figure 3.28 CAD drawings of the proposed passive shielding design for this experiment. The overall shielding system comprises several shield bodies with end caps (left) which are nested concentrically (right). This shielding design provides an extremely high shielding factor both radially and axially in simulations and is thus expected to be suitable for ultralow magnetic field experiments with ultracold atoms.

tens of cm) this leads to a relatively detailed mesh over large areas, giving rise to a great number of vertex elements for the relevant problem equations to be solved over. This naturally leads to lengthy computation times with a single iteration of computing the flux density distribution over the geometry taking multiple hours and the requisite number of iterations to achieve a sufficiently precise result taking multiple days.

The problem is posed as a spherical boundary with a known value of field (set to 50 μT, typical of Earth's field) containing our shields (set as having relative permeability $\mu_r = 5 \times 10^4$, typical of μ-metal). From the known, fixed field values at the boundaries of the problem (which are sufficiently far from the shields to avoid proximity effects) fields are iteratively calculated over the entire problem geometry in order to evaluate the flux density distribution within the shields. A minimum mesh size of 0.5 mm was used in order to capture edge effects at the mouth of the

shield's tubulations. The relative alignment of the external field with the axis of the shields was varied between parallel with the axis of the cylinders and perpendicular to this in order to capture the anisotropy of the shielding factor.

Results from these simulations, shown in figure 3.29 estimate that the axial shielding factor of our shield design is $S_{ax} \approx 10^{10}$ and the radial shielding factor is $S_{rad} \approx 10^{11}$. An error tolerance of 10^{-7} was used as a trade-off between accuracy and speed of simulation. This led to some numerical noise on the very lowest fields reached, as in the case of an external magnetic field perpendicular to the axis of the cylinders (see figure 3.29 (b)) where the shielding factor is very high. The exact details of this region of the simulation are not particularly interesting as the field is already well below our required target and the behaviour of the field is expected to be smooth in this region based on previous simulations. As expected, the $S_{ax} < S_{rad}$ however S_{ax} is still sufficiently higher than the required $S_{req} = 10^8$ to provide the low total field magnitude we require regardless of the alignment of the external field to perform our experiments. We have intentionally over-engineered the shielding performance (by a factor of around 100) in order that we have a large overhead wherein we can tolerate losing shielding performance, e.g. through manufacturing imperfections, mild mechanical shocks to the shields, imperfect degaussing, etc., without the integrity of the experiment being affected.

(a)

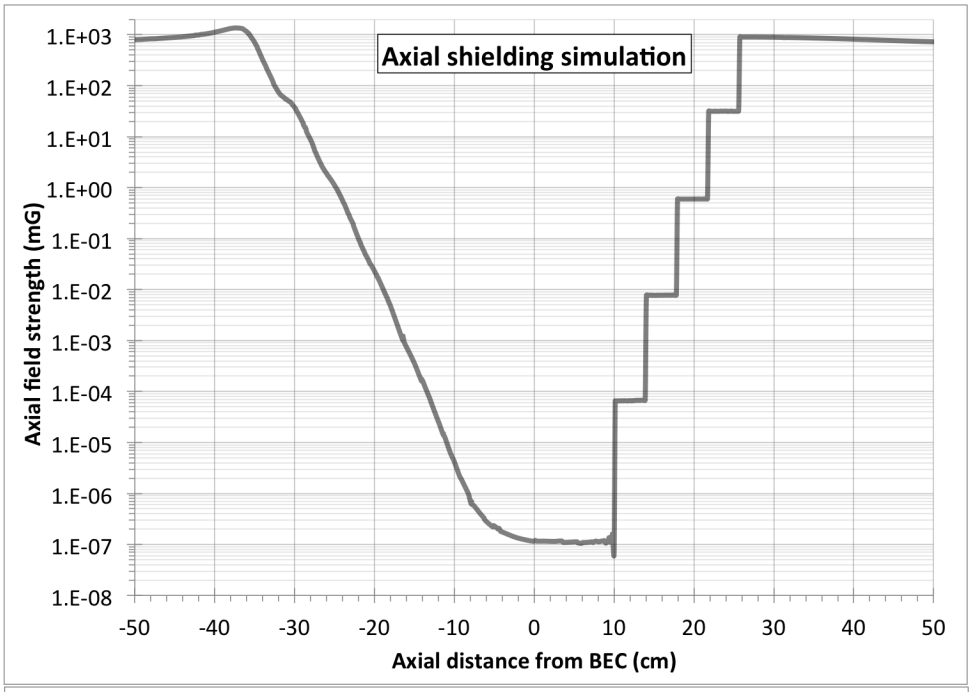

(b)

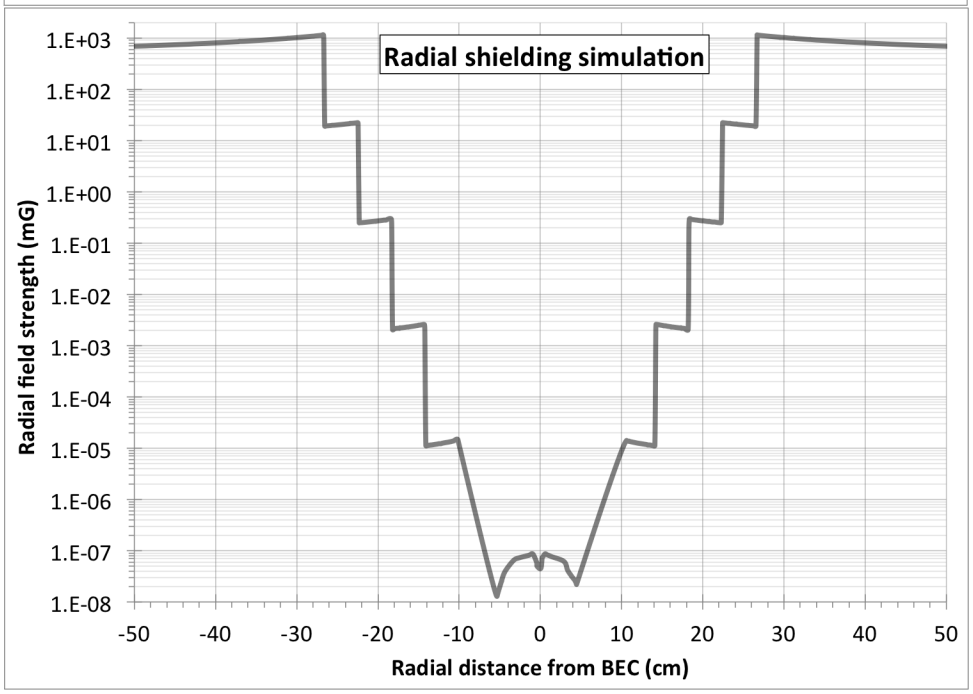

Figure 3.29 Results from simulations of our passive shielding design, shown in figure 3.28, for an external field of 50 μT aligned (a) with the shield axis and (b) orthogonal to the shield axis. The non-smooth noise seen in (b) is due to using an error tolerance of 1×10^{-7} for this simulation, meaning that changes below this level are viewed as converged. One can readily predict that the behaviour in this region will be a smooth change of field across the inside of the innermost shield to the values at the edges and do not require the results in this region. It is more important to us that these simulations tell us how the shielding performance varies along the axis of the shield. NB: we require fields of ≤ 1 μG ($= 1 \times 10^{-3}$ mG) in order to observe dipolar effects within our BEC, which these shields very comfortably achieve in simulation.

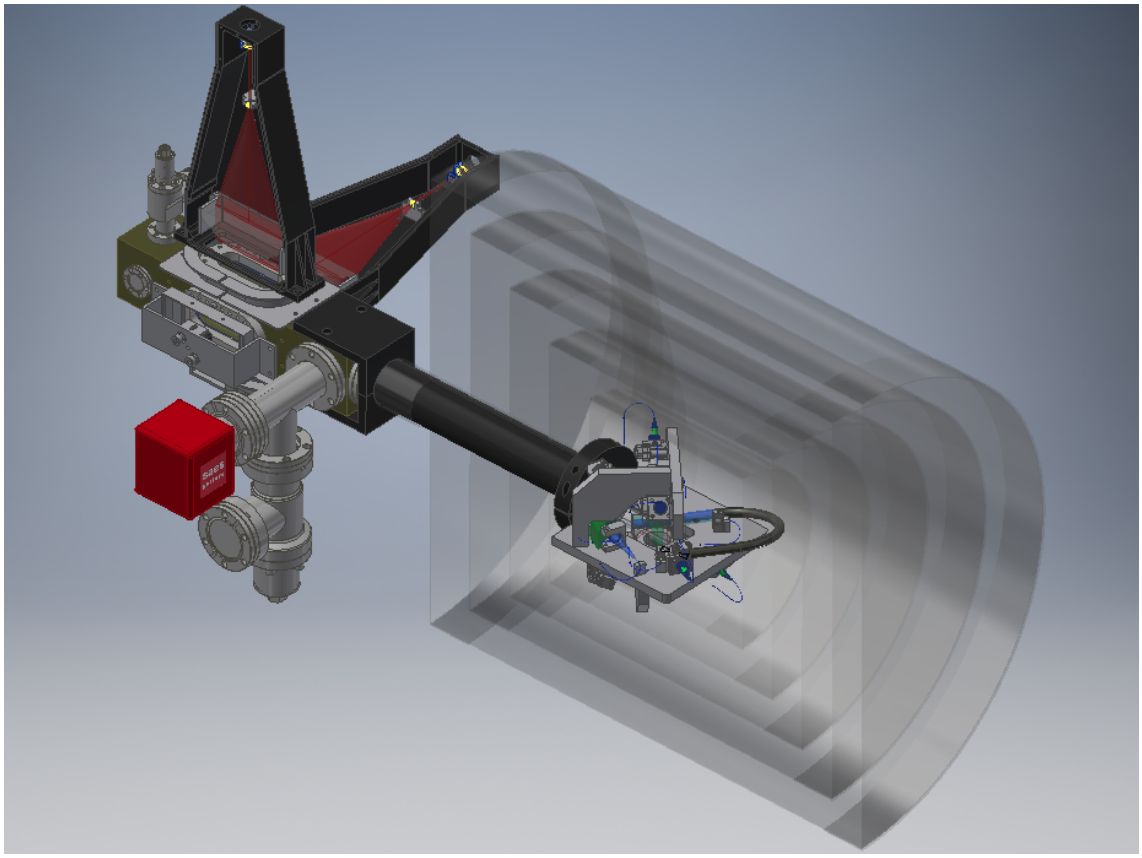

Figure 3.30 Our overall experimental setup. The main vacuum chamber (shaded green here) is seen on the left, with attached telescopes and other components to produce our 2D MOT. The glass cell, protruding to the right, is largely obstructed by mounts and optics for the 3D MOT, optical lattice, and dipole trap. Surrounding the glass cell optics are the mu-metal shields (shown here transparent and halved) employed to reach the ultralow fields required for our experiments.

4

Active magnetic shielding

Central to our experiment is the achievement and maintenance of an ultralow magnetic field environment. We have previously discussed the nature of the ambient magnetic field within the laboratory which can be separated into 4 parts: near-DC fields (<1 Hz), low-frequency AC fields (<1 kHz), high-frequency AC fields (>1 kHz), and transient or pulsed magnetic fields. The passive shielding system we have designed is highly effective at shielding both near-DC and >1 kHz AC fields however is less effective at shielding <1 kHz AC fields and transient fields. We know *a priori* that there will be significant and unavoidable noise sources in the region between DC and 1 kHz due to currents within the mains circuit (50 Hz and higher harmonics due to nonlinear loads), equipment with large motors (air conditioning

units, ca. 70 Hz, vacuum pumps, ca. >100 Hz), and from fluorescent lighting (ca. 100 Hz and higher harmonics due to using full-wave-rectification of the mains line). Further, one can reasonably expect to experience transient magnetic fields due to movement of large ferrous bodies such as cars on nearby roads, elevators, and magnetic equipment moved by hand. Active magnetic shielding systems can provide effective shielding from DC to their bandwidth (ca. 1 kHz) and importantly also against transient fields down to periods approaching their inverse bandwidth (ca. 1 ms). In order to protect against such noise affecting our experiment we chose to supplement our passive shielding system with an active one.

At their core, active compensation systems of any kind (e.g. for temperature[78], angular acceleration[79], or magnetic field[80–86]) comprise a sensing system coupled to a means of affecting the sensed quantity via a feedback loop. Being able to monitor the actual output, as opposed to the desired output, of the system enables the possibility of adjusting the desired output dynamically in order to stabilise the actual output. As an example, a fan with a fixed speed placed near an open window can be used to reduce the temperature of a room by increasing the flow of cool air, however this may cause the temperature of the room to become colder than desired when left on for a long period of time. On the other hand, an air conditioning unit is able to measure the room temperature and adjust its fan speed accordingly to achieve the desired room temperature indefinitely within some margin of error. This concept can be employed here to enable the output of a magnetic field sensor to drive an electromagnet via a Proportional-Integral-Derivative (PID) controller with the net effect being control (in a parametric sense) of the ambient magnetic field at the position of the sensor. The overall performance is dependent on the dynamic range, bandwidth and noise floor of the total system. In order to achieve the greatest performance one should design the feedback system to have a lower noise floor and higher bandwidth than (or similar to) the sensor used, such that the resulting system will be limited by the sensor. Several trade-offs have to be made

when designing such a system, for example between bandwidth and noise floor, and should be made with consideration of the performance requirements.

We are only able to effect active control of the ambient magnetic field within a region of space around the sensor, the size of which depends on the configuration of the coils used to produce the compensation field. We can calculate the expected size of this "controlled volume" by considering the profile of the magnetic field generated by our coils. These consist of two circular coils of copper wire (200 μm diameter, 10 turns) with an approximate diameter of 19.2 cm which are aligned axially and separated by approximately 10 cm. When currents are run through these coils in series, such that the direction of the magnetic field produced by each coil is the same, they produce a field profile along their common axis which depends on the ratio of the diameter, d, of the coils compared to the separation, l, along their common axis. For the Helmholtz condition of $l = d/2$ the field profile produced along the coils' common axis has a region around the midpoint of the two coils' centres which spatially satisfies the condition of $\frac{\partial^n}{\partial x^n}\mathbf{B}(x) = 0$ up to fourth order (i.e. $n = 4$)[87]. What this means practically is that the axial component of the magnetic flux density produced at the midpoint of the two coils has a very well-defined value over a certain region and any spatial field gradients introduced in this region can be deemed negligible.

Theory 4.1

A practical measure of a system's noise floor is the size of the peak-to-peak noise during normal operation. When this noise represents random fluctuations within the system, as when the system is minimally limited, this should be evenly distributed over all frequencies. Sampled data points of such noise follow a Gaussian distribution, the standard deviation of which can be estimated from the noise power spectral density and bandwidth of the system under test. The standard deviation

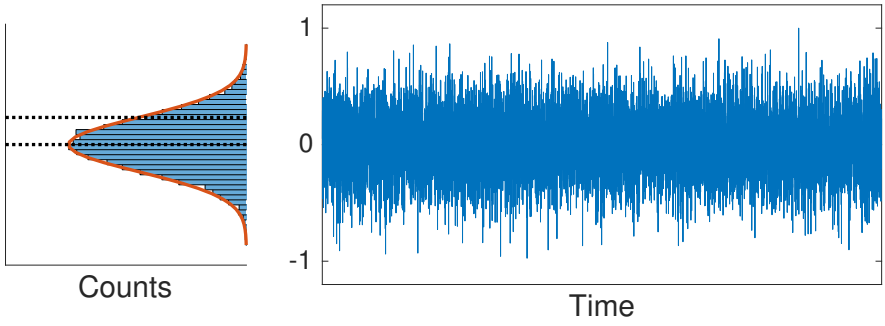

Figure 4.1 A sampled random signal (*right*) and a histogram of the data points, fitted to a Gaussian distribution (*left*). The black lines indicate one standard deviation (σ) of this distribution. For uncorrellated Gaussian noise, the peak-to-peak of the signal can be estimated as 6.6σ.

of noise can be calculated by multiplying the power spectral density (a single value over the spectrum for white noise) by the square root of the system's bandwidth. It follows that the noise within a system can be lowered by reducing the bandwidth, however one must consider the minimum bandwidth that can be tolerated before performance degrades due to a slower response.

The power spectral density (PSD) of a signal can be estimated using discrete timeseries data, such as that acquired from an oscilloscope. By comparing the PSD of the amplified sensor output with compensation on versus off we can see the effectiveness of the system over frequency. The power spectral density can be shown to be equivalent to the modulus squared of the discrete Fourier transform scaled by the frequency bin size via the Wiener-Khinchin theorem[88]. Taking many of such measurements and averaging them together allows us to lower the statistical error and provides a better picture of the true nature of the spectral distribution of noise within the system.

Design 4.2

The overall design of any electronic circuit, especially instrumentation, is somewhat iterative. A combination of parts need to be chosen initially and their performance

estimated then compared with the desired specification before modifying parts to better meet this specification. Fortunately, this circuit was mostly "right first time" with any remaining improvements offering little extra value compared to their required effort. This work was undertaken as part of a secondment within the group of Dr. Fabrice Gerbier at the Laboratoire Kastler Brossel of École Normale Supérieure. Some direction and guidance was given to me by Dr. Fabrice Gerbier and Dr. Tilman Zibold but the work is all my own. In the following sections I will present the components used, the expected performance from these parts, then compare this to the measured system performance.

Components 4.2.1

The compensation system used is based around an HMC1001 1-axis magnetic field sensor which delivers an exceptionally high bandwidth of 5 MHz. Internally the sensor comprises four anisotropic magnetoresistive elements wired in a Wheatstone bridge configuration, giving a differential output signal. Anisotropy of the resistive elements of the bridge gives rise to the directionality of the sensor and the differential nature of the output signal allows for common-mode noise rejection of electromagnetic interference picked up in the long lead wires connecting to the sensor. The sensor is powered with ± 5 V supplied by a buffered AD588 precision voltage reference to provide stable, low-noise power rails which can source a sufficiently high current to excite the internal resistive bridge of the sensor. The differential output of the sensor is amplified by a factor of 1000 using an AD524 precision instrumentation amplifier. Such an amplifier is chosen to provide the highest common-mode noise rejection on the sensor signal and to produce a high precision single-ended replication of the sensor signal for further processing. A PID controller is constructed using a number of OP07 op-amps which drive two voltage-controlled current sources (VCCS). One VCCS is a standard Howland current pump circuit [89] based around an op-amp (OPA633) which is able to drive a high current (using a BUF634 unity gain

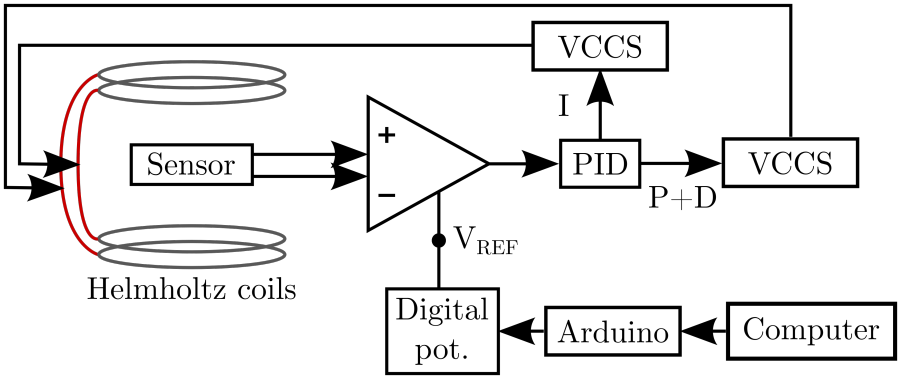

Figure 4.2 A schematic diagram of the acitve magnetic field compensation system described in this chapter. The upper half of the diagram is entirely analogue whereas the output offset of the sensor amplifier can be adjusted via a digital potentiometer, allowing us to interface this system with our experimental control computer.

buffer) and is designed to have a bipolar output with high bandwidth (> 100 kHz). A second VCCS is based on a low-side current sensing circuit using an op-amp (OPA2188) to servo a high power transistor (MJ11016) and is designed to drive very high currents with precision DC performance at the expense of having very low bandwidth (< 100 Hz). These two current sources are used as parallel feedback elements via two pairs of Helmholtz coils, to cancel DC and AC fields respectively. The I-output of our PID (i.e. low frequency feedback) is sent to the low bandwidth, high current VCCS to compensate for DC and very low frequency fields whilst the P- and D-outputs are sent to our high bandwidth, low current VCCS to compensate for AC fields. The use of these two separate feedback elements eases the requirements of the current sources. It was initially expected that there may be some interaction between the two pairs of coils due to their mutual inductance, however this was not observed in experiments. It is expected that this is because the coils themselves were tightly wound around each other with little space in between. A schematic diagram of the active compensation system can be seen in figure 4.2.

4.2.2 Bandwidth and noise

The bandwidth of the overall system inherently limits the maximum achievable compensation due to raising the noise floor by including higher frequencies of the white noise background arising from Johnson noise within the sensor. White noise has a uniform power spectrum which means that each 1 Hz bin contains the same amount of noise power, thus a higher system bandwidth includes more of these bins and in turn has a higher noise level. The Wheatstone bridge within the sensor has a sensitivity of 3.2 mV/V/G and in our case is powered with ±5 V = |10 V| giving an output of 32 mV/G. Thermally excited currents within the resistive elements of the bridge give rise to noise on the output voltage of the sensor. The resulting thermal voltage noise, V_{Nth}, due to thermal currents excited within a resistor, R, over a given bandwidth, Δf, at a temperature, T, can be calculated using[90]

$$V_{Nth} = \sqrt{4kTR\Delta f}, \tag{4.1}$$

where k is Boltzmann's constant. Noise figures are more commonly presented in terms of spectral noise density in units of $\mathrm{V}/\sqrt{\mathrm{Hz}}$ in a 1 Hz bandwidth as the noise density may change over the bandwidth of the device, which can be conveyed using a graph of the spectral noise density against frequency. For the HMC1001 bridge resistance of 850 Ω at 25°C a voltage noise spectral density of 3.74 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$ is produced by thermal currents, which is only slightly less than the datasheet quoted noise figure of 3.8 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$. The AD524 amplifying the sensor signal has an input-referenced voltage noise spectral density of 7 $\mathrm{nV}/\sqrt{\mathrm{Hz}}$, bringing the noise floor up to $\sqrt{3.8^2 + 7^2} \approx 8$ $\mathrm{nV}/\sqrt{\mathrm{Hz}}$. The AD524 has a gain bandwidth product (GBP) of 25 MHz and produces a gain of 1000, giving the amplified output signal a GBP-limited bandwidth of 25 kHz. It follows that the RMS voltage noise at the output

of the combined sensor and amplifier system is

$$V_{\mathrm{RMS}} = \left(8\ \mathrm{nV}/\sqrt{\mathrm{Hz}}\right)\left(\sqrt{25\ \mathrm{kHz}}\right)(1000) = 1.26\ \mathrm{mV}(\equiv 40\ \mu\mathrm{G}). \qquad (4.2)$$

As the RMS amplitude of a sampled white noise signal with no DC component is equal to the standard deviation (σ) of the samples, a peak-to-peak (pk-pk) amplitude for a defined percentage of the samples can be calculated by multiplying V_{RMS} by the number of standard deviations required to include this percentage of samples. The industry standard is to include 99.9% of samples, or equivalently 6.6 σ. Hence from our RMS noise amplitude we can calculate a pk-pk noise amplitude of 8349 nV $\equiv$ 260 μG. The residual noise observed at the output of the AD524 after tuning the system to provide optimal compensation was 12 mV $\equiv$ 375 μG, thus we expect that with this configuration we are approaching the electronic noise floor of the system. In order to achieve a pk-pk noise level of 100 μG by simply reducing the bandwidth we would need to lower the bandwidth to around 3.5 kHz, highlighting the diminishing returns from the square-root scaling of the noise level with bandwidth. Such a diminuiton of the bandwidth would certainly affect transient performance but this trade-off may be worthwhile as most steady state noise sources are be <1 kHz.

The datasheet for the HMC1001 mentions that the resolution of the sensor when powered with 5 V is 27 μG in a 10 Hz bandwidth. Assuming this bandwidth starts at DC, we will need to take the 1/f noise into account which is 30 nV/$\sqrt{\mathrm{Hz}}$ at 1 Hz. Repeating our previous calculations for the region 0.1–10 Hz we expect a pk-pk noise of $6.6(30\sqrt{\ln(10/0.1)}\ \mathrm{nV})/(32\ \mathrm{mV/G}) = 13\ \mu$G for this sensitivity and bandwidth. The difference between this and the value quoted in the datasheet is due to our powering the sensor with 10 V rather than 5 V. The resolution quoted is limited by 1/f noise near DC and thus field resolution should be higher at higher frequencies (i.e. smaller fields can be detected within the same bandwidth at higher frequencies where we only have white noise).

Offset 4.2.3

The datasheet for the HMC1001 states that the zero offset for a bridge voltage of 8 V can vary from −60 mV (−2.3 G) to 30 mV (1.2 G), with a typical offset of −15 mV (−0.59 G). The equivalent field figures are based on the average sensitivity of 3.2 mV/V/G, however this can also vary from unit to unit being as low as 2.5 mV/V/G meaning this offset could equivalently range from −3 G to 1.5 G in the worst case. In our case we are able to account for this offset by changing the output reference voltage of the amplifier. Residual DC fields will be highly attenuated by passive shielding around the experimental region, however such offsets should be considered if the absolute DC field after compensation is critical.

Performance 4.3

In practical compensation, 50 Hz from mains is the dominant component of the noise spectrum. It is easy to measure small signals at this frequency by line-triggering an oscilloscope and measuring the sensor signal with suitable averaging. Using this method, a residual 50 Hz pk-pk of 33.2 mV ≡ 1.04 mG was measured without any compensation and 1.8 mV ≡ 56 μG with compensation. This measurement was made with a sampling interval of 50 ms and thus a Fourier-limited bandwidth of 1/2(50 ms) = 10 Hz which should have comparable noise to the figures calculated in section 4.2.2. These results demonstrate field detection below the broadband noise level (260 μG) for a given frequency component and approaches the limits of the system's sensitivity based on the electronic noise floor (27 μG in a 10 Hz bandwidth).

The PSD of the sensor output with and without active compensation can be seen in figure 4.3. It can be seen that, as expected, the ambient noise power mainly consists of a peak at 50 Hz due to mains currents and higher harmonics of this. The effectiveness of the active compensation is clearly demonstrated by the near flat noise spectrum when the system is on, demonstrating a reduction of noise power

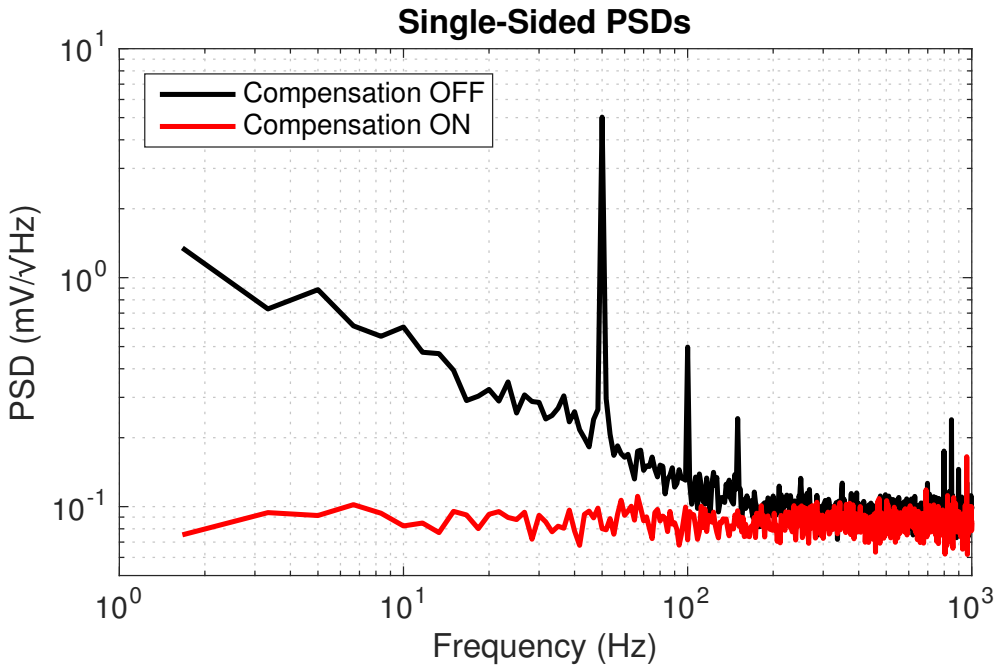

Figure 4.3 Power spectral densities of the amplified sensor output with (*red*) and without (*black*) active compensation. One can clearly see that without compensation applied there is a large amount of noise power at the 50 Hz mains frequency and higher harmonics of this. When active compensation is turned on the noise floor is brought down to that of the op-amps in the feedback loop, leaving no features in the spectrum.

at 50 Hz by more than 3 orders of magnitude. With compensation applied the noise floor is limited by that of the op-amps used to implement the feedback loop at around 8 nV/$\sqrt{\text{Hz}}$, however this is close to the limit of the sensor's intrinsic noise at around 4 nV/$\sqrt{\text{Hz}}$.

The bandwidth of the system was estimated by measuring the amplitude of an injected noise signal at a given frequency with the compensation system on versus off. This noise signal was generated by driving a low inductance coil with a sine wave from a frequency generator via a voltage-controlled current source with high bandwidth. Acquiring this data on an oscilloscope triggered on the noise signal it is possible to use averaging to reduce noise from other frequency components in order to get a more accurate measurement at each point. The magnetic field reduction with the compensation system turned on versus frequency can be seen in figure 4.4. It can be seen that the system essentially behaves as a low-pass filter and provides

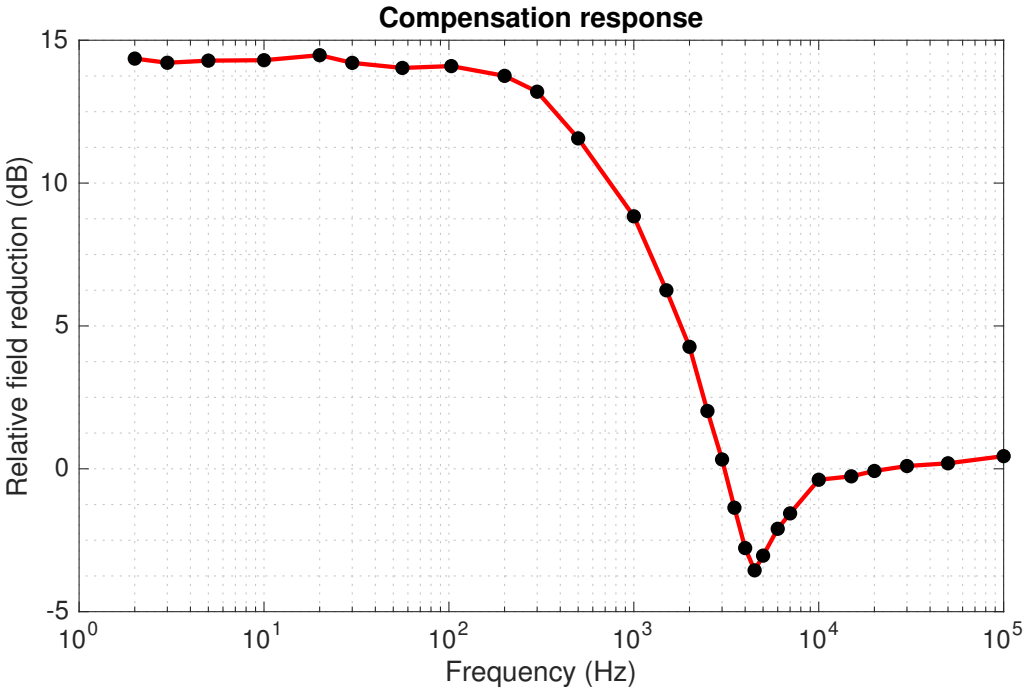

Figure 4.4 Reduction of noise versus frequency for the active compensation system described here. The black points are data and the red line is a guide to the eye. The amount of compensation is reduced by 10 dB with respect to low frequencies at around 2 kHz, which can be thought of as the bandwidth of the compensation system.

no reduction in noise at around 3 kHz, which we take to be the bandwidth of the compensation system.

As well as the steady-state performance of the compensation system, we are also interested in the transient performance as we may require to switch magnetic fields from one stable value to another within our experimental sequence. In order to measure the settling time of the system to a step change in field I again injected noise, with the compensation system on, in the same manner as above but with a square wave rather than a sine wave. The result of this measurement can be seen in figure 4.5. The step change of magnetic field causes the feedback loop to ring at its resonant frequency (typically near the bandwidth of the system) then decay back to its steady-state. Fitting this with an exponentially decaying cosine wave we can extract a decay time constant of $\tau = 145$ μs. For an exponential decay, one can assume that steady-state is reached again after roughly $5\tau = 725$ μs, meaning that

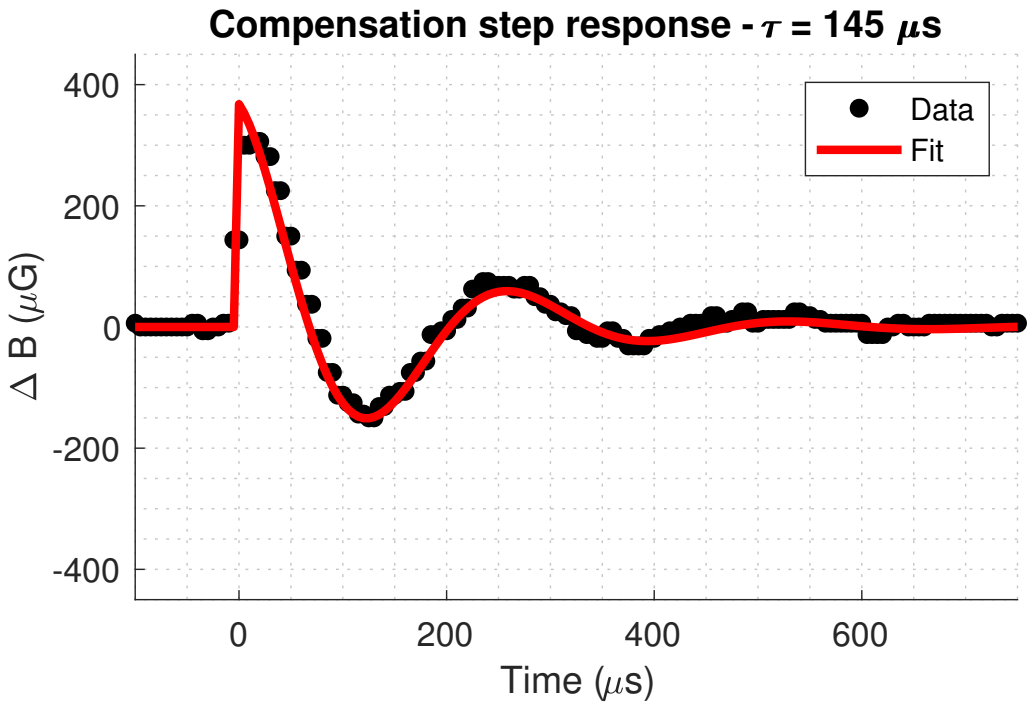

Figure 4.5 Output of the amplified sensor to a step change in field with active compensation applied. The data is fitted to an exponentially decaying cosine wave, giving rise to a decay time constant $\tau = 145$ μs. After 5τ has passed the output signal should have returned to its steady-state value, giving rise to a minimum response time of around 725 μs (i.e. < 1 ms), which is sufficiently rapid to be altered as part of an automated experimental sequence for many cold atoms experiments.

our compensation system is able to provide stable fields after a change in set-point or ambient field on a millisecond timescale. This is well-suited to most cold atoms experiments where much of the preparation stages of an experimental run take place on a millisecond timescale.

Based on the points discussed above we can assume that this system is optimised to near the electronic limits of the system for the bandwidth we have chosen. The noise floor could be reduced by using a different instrumentation amplifier with a lower voltage spectral noise density and/or reducing the system bandwidth. The former should be more effective and without detriment whereas the latter would affect transient performance, which is undesirable. As testament to the value and quality of this work this system has been replicated by collaborators at the Institut für Laserphysik of Universität Hamburg and is also intended to be installed in

collaborators' labs at the Laboratoire Kastler Brossel of École Normale Supérieure. A publication of the design, implementation, and performance of the active magnetic shielding system described in this thesis and its integration into an ultracold atoms experiment is currently being prepared with collaborators at the Institut für Laserphysik of Universität Hamburg and the Laboratoire Kastler Brossel of École Normale Supérieure.

5

Outlook

The results presented in this thesis have outlined the design and construction of an experimental apparatus capable of producing an ultralow magnetic field environment where an atomic physics experiment can be performed. This was achieved through careful material choices around the experiment region, design of a high performance passive magnetic shield, and development of a state-of-the-art active magnetic field compensation system. Several issues were encountered during the assembly and testing of the vacuum system which hindered the overall progress of the experiment, with the methods presented here being solutions to the problems faced. Specifically, issues with indium sealing of rectangular windows onto flat surfaces were particularly troublesome. The method developed using a milled channel to provide a sealing

face was very effective and solves several issues regarding indium sealing in general. Further issues regarding potential microleaks within the epoxy seals of our glass cell caused delays to progress but seem to have been resolved by switching to a frit glass bonded cell.

Potential future issues for the experiment are that stray reflections by beams in the glass cell due to the uncoated internal faces may cause interference fringes in our images. These can be mitigated using Fourier filtering of the images to selectively remove the spatial frequency of these fringes. Further concerns are the long-term stability of the optical alignment of components within the region around the glass cell. With the passive shielding in place, these components are not directly accessible and as such are designed to be "set and forget". It is likely that some re-alignment will be necessary after a period of time however if this is required on a weekly or monthly basis then this could be detrimental as each removal and replacing of the passive shield's end caps risks damaging the shield. Whilst the performance of our passive shielding has been simulated it is unclear how well this will perform in practice. To mitigate the effects of manufacturing imperfections and mechanical shocks to the shield the performance has been specified two orders of magnitude higher than is required. It is hoped that the performance of our active compensation system will be able to ease the requirements of the passive shield further by reducing and stabilising the ambient magnetic field in the volume which the passive shields sit.

Despite the technical challenges ahead, a significant number of major issues regarding the realisation of this experiment have been overcome. The apparatus is currently operational (without passive shielding) and we have been able to demonstrate loading of the 3D MOT. From here we need only align our dipole trap beam with the 3D MOT and begin optimisation of our experimental sequence (regarding ramps and timings) in order to load the dipole trap and attempt to evaporatively cool to BEC. This will be a large milestone in the project, proving that the under-

lying apparatus is fit for purpose. From then the next challenge will be to test and integrate the passive shielding, ensuring that BEC can still be achieved inside of the shields. Once this is done, we will have realised our ambitions of producing BEC in an ultralow magnetic field and will be poised to perform novel experiments in this previously unreached parameter regime.

With the wealth of avenues of research available with such a system, including testing the Einstein-de Haas effect in a quantum system[17–20], producing macroscopic entanglement[29,30], and observing novel ground state phases[14–16], our main question will be — what to do first?

References

[1] Plato. *Ion*. Urbana, Illinois: Project Gutenberg, 2008.

[2] C. Bailey. *Lucretius: On the Nature of Things*. Clarendon Press, 1946.

[3] C. Plinius Secundus. *Pliny's Natural History in Thirty-seven Books (Translated by Dr. Philemon Holland)*. The Wernerian Club, 1847.

[4] J. B. Porta. *Natural Magick: in* XX *bookes*. 1658.

[5] W. Gilbert. *On The Loadstone and Magnetic Bodies and on The Great Magnet The Earth. A New Physiology, Demonstrated with many arguments and experiments (Translated by P. Fleury Mottelay)*. Ferris Brothers, 1893.

[6] K. J. Arnold and M. D. Barrett. All-optical Bose-Einstein condensation in a 1.06 μm dipole trap. *Opt. Comm.*, 284(13), 2011.

[7] N. E. Lundblad. *All-optical spinor Bose-Einstein condensation and the spinor dynamics-driven atom laser.* PhD thesis, California Insitute of Technology, 2006.

[8] M. Barrett, J. Sauer, and M. Chapman. All-Optical Formation of an Atomic Bose-Einstein Condensate. *Phys. Rev. Lett.*, 87(1), 2001.

[9] T. Kinoshita, T. Wenger, and D. Weiss. All-optical Bose-Einstein condensation using a compressible crossed dipole trap. *Phys. Rev. A*, 71(1), 2005.

[10] S-J. Deng, P-P. Diao, Q-L. Yu, and H-B. Wu. All-Optical Production of Quantum Degeneracy and Molecular Bose-Einstein Condensation of ^{6}Li. *Chinese Phys. Lett.*, 32(5), 2015.

[11] J. A. Hughes. *An all optical BEC apparatus for the investigation of magnetic order in spinor atomic gases in ultralow magnetic fields.* PhD thesis, University of Georgia, 2013.

[12] N. E. Lundblad, R. J. Thompson, D. C. Aveline, and L. Maleki. Initial experiments with an all-optical spinor BEC. In *Quant. Elec. Laser Sci. Conf.* OSA, 2006.

[13] D. A. Steck. Rubidium 87 D Line Data. *J. Geophys. Res.*, 2009(2), 2010.

[14] S. Yi and H. Pu. Spontaneous Spin Textures in Dipolar Spinor Condensates. *Phys. Rev. Lett.*, 97(2), 2006.

[15] M. Ueda. Bose Gases with Non-zero Spin. *Ann. Rev. Cond. Mat. Phys.*, 3(1), 2012.

[16] S. Yi and H. Pu. Dipolar spinor Bose-Einstein condensates. *arXiv*, 0804.0191(1), 2008.

[17] K. Gawryluk, M. Brewczyk, K. Bongs, and M. Gajda. Resonant Einstein-de Haas Effect in a Rubidium Condensate. *Phys. Rev. Lett.*, 99(13), 2007.

[18] M Ganzhorn, S. Klyatskaya, M. Ruben, and W. Wernsdorfer. Quantum Einstein-de Haas effect. *Nat. Commun.*, 7(11443), 2016.

[19] Y. Kawaguchi, H. Saito, and M. Ueda. Einstein-de Haas Effect in Dipolar Bose-Einstein Condensates. *Phys. Rev. Lett.*, 96(8), 2006.

[20] T. Świsłocki, T. Sowiński, J. Pietraszewicz, M. Brewczyk, M. Lewenstein, J. Zakrzewski, and M. Gajda. Tunable dipolar resonances and Einstein-de Haas effect in a ^{87}Rb-atom condensate. *Phys. Rev. A*, 83(063617), 2011.

[21] G. M. Wysin, W. A. Moura-Melo, L. A. S. Mól, and A. R. Pereira. Dynamics and hysteresis in square lattice artificial spin ice. *New J. Phys.*, 15(4), 2013.

[22] A. Remhof, A. Schumann, A. Westphalen, H. Zabel, N. Mikuszeit, E. Y. Vedmedenko, T. Last, and U. Kunze. Magnetostatic interactions on a square lattice. *Phys. Rev. B*, 77(13), 2008.

[23] R. F. Wang, C. Nisoli, R. S. Freitas, J. Li, W. McConville, B. J. Cooley, M. S. Lund, N. Samarth, C. Leighton, V. H. Crespi, and P. Schiffer. Artificial 'spin ice' in a geometrically frustrated lattice of nanoscale ferromagnetic islands. *Nature*, 439(7074), 2006.

[24] E. Y. Vedmedenko and N. Mikuszeit. Multipolar ordering in electro- and magnetostatic coupled nanosystems. *ChemPhysChem*, 9(9), 2008.

[25] G-W. Chern and O. Tchernyshyov. Magnetic charge and ordering in kagome spin ice. *Philos. Trans. R. Soc. A*, 370, 2012.

[26] R. V. Hügli, G. Duff, B. O'Conchuir, E. Mengotti, L. J. Heyderman, A. Fraile Rodriguez, F. Nolting, and H. B. Braun. Emergent magnetic monopoles, dis-

order, and avalanches in artificial kagome spin ice (invited). *J. Appl. Phys.*, 111(7), 2012.

[27] S. T. Bramwell, S. R. Giblin, S. Calder, R. Aldus, D. Prabhakaran, and T. Fennell. Measurement of the charge and current of magnetic monopoles in spin ice. *Nature*, 461(7266), 2009.

[28] E. Mengotti, L. J. Heyderman, A. Fraile Rodríguez, F. Nolting, R. V. Hügli, and H. B. Braun. Real-space observation of emergent magnetic monopoles and associated Dirac strings in artificial kagome spin ice. *Nat. Phys.*, 7(1), 2010.

[29] S. Yi and H. Pu. Magnetization, squeezing, and entanglement in dipolar spin-1 condensates. *Phys. Rev. A*, 73(2), 2006.

[30] C-L. Zang, J. Yu, W-L. Yang, M. Feng, and S-P. Kou. Preparation of NOON State Induced by Macroscopic Quantum Tunneling in an Ising Chain. *arXiv*, 1210.6754(1), 2012.

[31] W. Wernsdorfer and R. Sessoli. Quantum Phase Interference and Parity Effects in Magnetic Molecular Clusters. *Science*, 284(5411), 1999.

[32] I. K. Kominis, T. W. Kornack, J. C. Allred, and M. V. Romalis. A subfemtotesla multi-channel atomic magnetometer. *Nature*, 422(6932), 2003.

[33] H. B. Dang, A. C. Maloof, and M. V. Romalis. Ultrahigh sensitivity magnetic field and magnetization measurements with an atomic magnetometer. *Appl. Phys. Lett.*, 97(151110), 2010.

[34] M. J. P. Gingras. Observing monopoles in a magnetic analog of ice. *Science*, 326(5951), 2009.

[35] P. Milde, D. Köhler, J. Seidel, L. M. Eng, A. Bauer, A. Chacon, J. Kindervater, S. Mühlbauer, C. Pfleiderer, S. Buhrandt, C. Schütte, and A. Rosch. Unwinding of a skyrmion lattice by magnetic monopoles. *Science*, 340(6136), 2013.

[36] M. Ueda and Y. Kawaguchi. Skyrmions in spinor Bose-Einstein condensates. *arXiv*, 0002375(1), 2000.

[37] C. J. Foot. *Atomic Physics*. Oxford University Press, Oxford, 8th edition, 2011.

[38] H. J. Metcalf and P. Van der Straten. *Laser cooling and trapping*. Springer, 1999.

[39] G. Wasik and R. Grimm. Non-magnetic atom trap based on a 3D bichromatic optical superlattice. *Opt. Comm.*, 137, 1997.

[40] S. Yi, L. You, and H. Pu. Quantum Phases of Dipolar Spinor Condensates. *Phys. Rev. Lett.*, 93(4), 2004.

[41] K. Gross, C. Search, H. Pu, W. Zhang, and P. Meystre. Magnetism in a lattice of spinor Bose-Einstein condensates. *Phys. Rev. A*, 66(3), 2002.

[42] H. Pu, W. Zhang, and P. Meystre. Ferromagnetism in a Lattice of Bose-Einstein Condensates. *Phys. Rev. Lett.*, 87(14), 2001.

[43] B. Pasquiou, E. Maréchal, G. Bismut, P. Pedri, L. Vernac, O. Gorceix, and B. Laburthe-Tolra. Spontaneous Demagnetization of a Dipolar Spinor Bose Gas in an Ultralow Magnetic Field. *Phys. Rev. Lett.*, 106(25), 2011.

[44] A. de Paz, A. Chotia, E. Maréchal, P. Pedri, L. Vernac, O. Gorceix, and B. Laburthe-Tolra. Resonant demagnetization of a dipolar Bose-Einstein condensate in a three-dimensional optical lattice. *Phys. Rev. A*, 87(5), 2013.

[45] M. Fattori, T. Koch, S. Goetz, A. Griesmaier, S. Hensler, J. Stuhler, and T. Pfau. Demagnetization cooling of a gas. *Nat. Phys.*, 2(11), 2006.

[46] SAES getters. *Alkali Metal Dispensers*. Technical report, 2007. Available online. Accessed: 1st May 2017.

[47] J. Schoser, A. Batär, R. Löw, V. Schweikhard, A. Grabowski, Y. Ovchinnikov, and T. Pfau. Intense source of cold Rb atoms from a pure two-dimensional magneto-optical trap. *Phys. Rev. A*, 66(2), 2002.

[48] G. Humpston and D. M. Jacobson. Chapters 1 and 2. In *Principles of Soldering*. ASM International, 2004.

[49] R. G. Budynas and K. J. Nisbett. *Shigley's Mechanical Engineering Design*. McGraw-Hill, 2011.

[50] G. A. Parks and S. Akhtar. Magnetic moment of Fe^{2+} in paramagnetic minerals. *Am. Mineral.*, 53(1), 1968.

[51] W. H. Dennen. Impurities in quartz. *Geol. Soc. Am. Bull.*, 75, 1964.

[52] ShinEtsu Quartz Products Co. Ltd. and Heraeus. *Quartz Glass for Optics*. Technical report, 2010.

[53] G. Ritt, G. Cennini, C. Geckeler, and M. Weitz. Laser frequency offset locking using a side of filter technique. *Appl. Phys. B*, 79(3), 2004.

[54] J. M. Lockhart. SQUID readout and ultra-low magnetic fields for Gravity Probe-B (GP-B). *Cryo. Opt. Syst. Inst. II*, 619, 1986.

[55] A. W. Rücker. On the magnetic shielding of concentric spherical shells. *Proc. Phys. Soc. London*, 12(462), 1892.

[56] F. Schweizer. Magnetic Shielding Factors of a System of Concentric Spherical Shells. *J. Appl. Phys.*, 33(3), 1962.

[57] T. E. Sterne. Multi-Lamellar Cylindrical Magnetic Shields. *Rev. Sci. Inst.*, 6(10), 1935.

[58] W. G. Wadey. Magnetic Shielding with Multiple Cylindrical Shells. *Rev. Sci. Inst.*, 27(11), 1956.

[59] T. J. Sumner, J. M. Pendlebury, and K. F. Smith. Convectional magnetic shielding. *J. Phys. D*, 20(9), 2000.

[60] E. Paperno, M. V. Romalis, and Y. Noam. Optimization of five-shell axial magnetic shields having openings in the end-caps. *IEEE Trans. Magn.*, 40(4), 2004.

[61] A. J. Mager. Magnetic shields. *IEEE Trans. Magn.*, 6(1), 1970.

[62] A. K. Thomas. Magnetic shielded enclosure design in the DC and VLF region. *IEEE Trans. Electromag. Compat.*, 10(1), 1968.

[63] V. Kelha, J. Pukki, R. Peltonen, A. Penttinen, R. Ilmoniemi, and J. Heino. Design, construction, and performance of a large-volume magnetic shield. *IEEE Trans. Magn.*, 18(1), 1982.

[64] V. V. Yashchuk, S.-K. Lee, and E. Paperno. Magnetic shielding. In *Optical Magnetometry*. 2012.

[65] A. Schnabel, J. Voigt, M. Burghoff, S. Kanppe-Gruneberg, C. Kittel, and F. Thiel. Magnetic Shielding. Technical report. Available online. Accessed: 1st May 2017.

[66] E. A. Donley, E. Hodby, L. Hollberg, and J. Kitching. Demonstration of high-performance compact magnetic shields for chip-scale atomic devices. *Rev. Sci. Inst.*, 78(8), 2007.

[67] T. Rikitake. Magnetic shielding by a spherical shell having a hole (1). *J. Geomag. Geoelec.*, 42, 1990.

[68] E. A. Burt and C. R. Ekstrom. Optimal three-layer cylindrical magnetic shield sets for scientific applications. *Rev. Sci. Inst.*, 73(7), 2002.

[69] V. Yashchuk, D. Budker, and M. Zolotorev. Applications of nonlinear magneto-optic effects with ultra-narrow widths. *AIP Conf. Proc.*, 457(177), 1999.

[70] R. M. Bozorth. *Ferromagnetism.* IEEE Press, New Jersey, 1st edition, 1993.

[71] T. W. Kornack, S. J. Smullin, S.-K. Lee, and M. V. Romalis. A low-noise ferrite magnetic shield. *Appl. Phys. Lett.*, 90(22), 2007.

[72] S. Malkowski, R. Adhikari, B. Hona, C. Mattie, D. Woods, H. Yan, and B. Plaster. Technique for high axial shielding factor performance of large-scale, thin, open-ended, cylindrical Metglas magnetic shields. *Rev. Sci. Inst.*, 82(7), 2011.

[73] S. Dickerson, J. M. Hogan, D. M. S. Johnson, T. Kovachy, A. Sugarbaker, S.-W. Chiow, and M. A. Kasevich. A high-performance magnetic shield with large length-to-diameter ratio. *Rev. Sci. Inst.*, 83(065108), 2012.

[74] J. Nenonen, J. Montonen, and T. Katila. Thermal noise in biomagnetic measurements. *Rev. Sci. Inst.*, 67(6), 1996.

[75] T. Varpula and T. Poutanen. Magnetic Field Fluctuations Arising From Thermal Motion of Electric Charge in Conductors. *J. Appl. Phys.*, 55(11), 1984.

[76] J. C. Allred, R. N. Lyman, T. W. Kornack, and M. V. Romalis. High-sensitivity atomic magnetometer unaffected by spin-exchange relaxation. *Phys. Rev. Lett.*, 89(13), 2002.

[77] J. F. Hoburg. Principles of quasistatic magnetic shielding with cylindrical and spherical shields. *IEEE Trans. Electromag. Compat.*, 37(4), 1995.

[78] S. Häfner, S. Falke, C. Grebing, S. Vogt, T. Legero, M. Merimaa, C. Lisdat, and U. Sterr. 8×10^{-17} fractional laser frequency instability with a long room-temperature cavity. *Opt. Lett.*, 40(9), 2015.

[79] M. Gajamohan, M. Merz, I. Thommen, and R. D'Andrea. The Cubli: A cube that can jump up and balance. In *IEEE Int. Conf. Intell. Robot. Syst.*, 2012.

[80] S. Celozzi and F. Garzia. Active shielding for power-frequency magnetic field reduction using genetic algorithms optimisation. *IEEE Proc. Sci. Meas. Tech.*, 151(1), 2004.

[81] D. Spemann, T. Reinert, J. Vogt, J. Wassermann, and T. Butz. Active compensation of stray magnetic fields at LIPSION. *Nucl. Inst. Meth. Phys. Res. B*, 210, 2003.

[82] C. J. Dedman, R. G. Dall, L. J. Byron, and A. G. Truscott. Active cancellation of stray magnetic fields in a Bose-Einstein condensation experiment. *Rev. Sci. Inst.*, 78(024703), 2007.

[83] M. Skakala, V. Zrubec, and J. Manka. Active compensation for ambient magnetic noise in the unshielded environment. *Meas. Sci. Tech.*, 4, 1993.

[84] L. A. Marzetta. Use of an Operational Amplifier with Helmholtz Coils for Reducing ac Induced Magnetic Fields. *Rev. Sci. Inst.*, 32(11), 1961.

[85] S. Kuriki, A. Hayashi, T. Washio, and M. Fujita. Active compensation in combination with weak passive shielding for magnetocardiographic measurements. *Rev. Sci. Inst.*, 73(2), 2002.

[86] D. Platzek, H. Nowak, F. Giessler, J. Röther, and M. Eiselt. Active shielding to reduce low frequency disturbances in direct current near biomagnetic measurements. *Rev. Sci. Inst.*, 70(5), 1999.

[87] David J. Griffiths. *Introduction to Electrodynamics*. Pearson, 3rd edition, 1999.

[88] R. N. Bracewell. *The Fourier transform and its applications*. McGraw-Hill, 3rd edition, 1999.

[89] Jacob Fraden. *Handbook of Modern Sensors: Physics, designs, and applications*. Springer, 3rd edition, 2003.

[90] H. Nyquist. Thermal agitation of electric charge in conductors. *Phys. Rev.*, 32, 1928.